

Cosmology at high redshift - a probe of fundamental physics

Noah Sailer,^{a,b} Emanuele Castorina,^{d,e} Simone Ferraro^{c,b} and Martin White^{a,b,c}

^aDepartment of Physics, University of California, Berkeley, CA 94720, USA

^bBerkeley Center for Cosmological Physics, UC Berkeley, CA 94720, USA

^cLawrence Berkeley National Laboratory, One Cyclotron Road, Berkeley, CA 94720, USA

^dDipartimento di Fisica ‘Aldo Pontremoli’, Università degli Studi di Milano, Milan, Italy

^eTheoretical Physics Department, CERN, 1211 Geneva 23, Switzerland

E-mail: nsailer@berkeley.edu, emanuele.castorina@unimi.it, sferraro@lbl.gov, mwhite@berkeley.edu

Abstract. An observational program focused on the high redshift ($2 < z < 6$) Universe has the opportunity to dramatically improve over upcoming LSS and CMB surveys on measurements of both the standard cosmological model and its extensions. Using a Fisher matrix formalism that builds upon recent advances in Lagrangian perturbation theory, we forecast constraints for future spectroscopic and 21-cm surveys on the standard cosmological model, curvature, neutrino mass, relativistic species, primordial features, primordial non-Gaussianity, dynamical dark energy, and gravitational slip. We compare these constraints with those achievable by current or near-future surveys such as DESI and Euclid, all under the same forecasting formalism, and compare our formalism with traditional linear methods. Our Python code [FishLSS](#) – used to calculate the Fisher information of the full shape power spectrum, CMB lensing, the cross-correlation of CMB lensing with galaxies, and combinations thereof – is publicly available.

Contents

1	Introduction	2
2	Surveys and probes	2
2.1	Near-term galaxy surveys	4
2.2	Future galaxy surveys	6
2.3	21-cm intensity mapping	8
2.4	CMB experiments	8
3	Forecasting formalism	9
3.1	Full shape measurements	10
3.2	CMB lensing and cross-correlations	13
3.3	Combined forecasts	15
3.4	Derivatives	16
3.5	Alcock-Paczynski effect	17
3.6	BAO modeling	18
3.7	Redshift uncertainties	18
4	Results	19
4.1	Distance measures	20
4.2	Base Λ CDM	21
4.3	Structure growth	23
4.4	Massive neutrinos	25
4.5	Light relics	27
4.6	Primordial non-Gaussianity	29
4.7	Primordial features	31
4.8	Dynamical dark energy	32
4.9	Modified gravity	35
4.10	Extended base model and multi-extension fits	36
4.11	Linear vs. non-linear modeling	37
5	Observational challenges	38
6	Conclusions	41
A	Primordial figure of merit (N_{modes})	43
B	Covariance approximations	45
C	BAO forecasting	46
D	Lensing auto- vs. cross-correlation	47
E	Empirical checks	49
F	Python implementation	50

1 Introduction

Large Scale Structure (LSS), as probed by fluctuations in the Cosmic Microwave Background (CMB) radiation or large galaxy surveys, provides one of our best windows on fundamental physics. Topics include, but are not limited to, constraints on the expansion history and curvature [1], primordial non-Gaussianity [2], features in the power spectrum (primordial [3] or induced [4–8]) or running of the spectral index [9, 10], models of inflation [11], dark energy [12, 13], dark matter interactions [14], light relics and neutrino mass [15, 16] and modified gravity [1, 17–19]. Unfortunately, while many of these extensions of the standard model are well motivated, we do not have strong guidance from theory on which of these many probes is most likely to turn up evidence of new, beyond standard model, physics. However, well understood phenomenology allows us to forecast where our sensitivities to such physics will be highest and our inference cleanest. To work where the inference is cleanest and the noise lowest we should push to high redshift.

Continuous advances in detector technology and experimental techniques are pushing us into a new regime, enabling mapping of large-scale structure in the redshift window $2 < z < 6$ using both relativistic and non-relativistic tracers. This will allow us to probe the metric, particle content and *both* epochs of accelerated expansion (Inflation and Dark Energy domination) with high precision. Moving to higher redshift allows us to take advantage of four simultaneous trends. First, a wider lever arm in redshift leads to rotating degeneracy directions among cosmological parameters and thus tighter constraints. Second, the volume on the past lightcone increases dramatically, leading to much tighter constraints on sample-variance dominated modes. Third, the degree of non-linearity is smaller, so the field is better correlated with the early Universe. Fourth, high precision perturbative models built around principles familiar from high-energy particle physics become increasingly applicable. Indeed at high redshift, we increasingly probe long wavelength modes which are linear or quasi-linear and carry a great deal of “unprocessed” information from the early Universe.

In this paper we forecast the constraints on cosmological parameters and extensions to the standard model that could be achieved by an observational program focused on the high redshift Universe. We compare the reach of different experiments using a Fisher matrix formalism that builds upon recent advances in cosmological perturbation theory. Our Python code `FishLSS`¹ used to create these forecasts is publicly available, and is described in greater detail in Appendix F.

Since at high redshift the non-linear scale shifts to higher k , perturbative models provide a robust method for forecasting that allows for consideration of scale dependent bias, stochastic small-scale power and the decorrelation of the initial and observed modes (see also ref. [20]). We begin in §2, where we list observational probes and discuss our assumptions about their number density and linear bias, and describe the experimental parameters of the (future) surveys for which we forecast. Details of the forecasting formalism we employ are provided in §3 and our results are presented in §4. We explore how these results depend on observational limitations in §5. We conclude in §6.

2 Surveys and probes

In this section we describe the various surveys and probes considered in our forecasts, as summarized in Table 1. These include spectroscopic galaxy surveys, which observe some

¹<https://github.com/NoahSailer/FishLSS>

combination of emission line galaxies (ELGs), H α galaxies, or Lyman break galaxies (LBGs), along with interferometers dedicated to 21-cm intensity mapping, and CMB observatories.




Figure 1. $b^2 n$ for each of the LSS surveys considered in our forecasts. This combination of linear bias and number density is proportional to the signal-to-noise ratio in linear theory, at fixed \mathbf{k} (§3). For both PUMA and HIRAX, the “number density” is taken to be the inverse of the thermal + shot noise. We note that while the fiducial PUMA design would enable observations down to $z = 0.2$ (dotted lines), we neglect this region in our forecasts, as it will be well measured by upcoming surveys such as DESI or Euclid which we already include.

In the following subsections we discuss our assumptions about the linear bias and number density of each LSS probe. These values are summarized in Fig. 1, where we plot the effective number density $b^2 n$ of each survey. This figure already provides some baseline intuition for the results of the following forecasts. In particular, we see that MegaMapper has a comparable noise to the Euclid H α sample, but covers a much broader redshift range, surveying a volume ~ 2.5 times as large, as illustrated in Fig. 4. Similarly, PUMA (-5K and -32K) has a lower noise than either MegaMapper, Euclid or DESI, and its survey volume is ~ 3.2 times that of Euclid². In addition to lower shot noise and more volume, high redshift ($z > 2$) experiments have access to more “linear” modes, i.e. modes that can be modeled very accurately using perturbation theory. In Fig. 2 we plot the “effective number of linear (primordial) modes” (N_{modes}) that MegaMapper, PUMA, and two upcoming CMB observatories will have access to. For a precise definition of N_{modes} , as well as a discussion of its interpretation and caveats on the comparison between CMB and LSS experiments using this metric, see Appendix A. In particular, we see that MegaMapper will have access to roughly 25% more modes than SO, and that PUMA-32K will have access to $\gtrsim 4$ times more modes than CMB-S4, assuming optimistic foregrounds. Our definition of N_{modes} can be easily modified to forecast the error on a measurement of the linear matter power spectrum $P_m(k, z = 0)$, as shown in Fig. 3. In particular we see that MegaMapper and PUMA will be capable of measuring the power spectrum out to $k = 0.9 h \text{Mpc}^{-1}$ with an order of magnitude higher precision than Planck at $k \sim 0.2 h \text{Mpc}^{-1}$, despite having finer-grained k -bins. These estimates suggest that cosmological constraints from future, high-redshift LSS surveys will be

²This rough factor ignores differences in sky coverage.

competitive with CMB measurements. As we show in §4, combining the two dramatically improves constraints on Λ CDM parameters, N_{eff} , the neutrino masses and other extensions of the standard cosmological model.




Figure 2. N_{modes} as a function of z_{max} for PUMA (-5K and -32K) and MegaMapper ($2 < z < z_{\text{max}}$). For PUMA, we consider both optimistic and pessimistic foreground models (§3.1), which are the boundaries of the shaded regions. In green are the number of modes for Simons Observatory and CMB-S4 like experiments with $7 \mu\text{K-arcmin}$ and $1.5 \mu\text{K-arcmin}$ noise levels respectively, both with a $1.4'$ beam and $f_{\text{sky}} = 0.4$.

We note that the forecasts presented in this work do not include constraints on cosmological parameters coming from cosmic shear or galaxy-galaxy lensing. Both quantities will be measured to exquisite precision by upcoming surveys such as Euclid, the Rubin Observatory and the Roman Space Telescope. However, most of their cosmological constraints from lensing will come from $z < 2$, due to the difficulty in defining high number density samples with well-characterized photometric redshifts and shape measurements at high redshift. Moreover, due to the relatively large redshift uncertainties compared to a spectroscopic survey, detailed features in the power spectrum (such as the BAO) are suppressed, together with the majority of the modes in the radial direction. Therefore, the analysis is often performed in a small number (order 10) of 2-dimensional tomographic bins. For this reason, photometric surveys are very effective at measuring the amplitude of structure and/or lensing at low redshift, while spectroscopic surveys can measure the 3-dimensional power spectrum and hence obtain growth and expansion measurements directly. Thus the high-redshift spectroscopic surveys studied here are highly complementary to other upcoming and planned photometric experiments.

2.1 Near-term galaxy surveys

The next five years will see a massive influx of LSS data, with spectroscopic galaxy surveys including DESI, Euclid, and SPHEREx coming online. Together these surveys will map out the low redshift universe, taking spectra of hundreds of millions of galaxies.

The Dark Energy Spectroscopic Instrument (DESI) [23] is actively taking data on the 4 m Mayall Telescope. Using its five thousand fibers and 10 spectrographs, DESI will take




Figure 3. Measurements of the linear matter power spectrum at $z = 0$. For both MegaMapper and PUMA-32K we show constraints for 15 linearly space k -bins between $0.1 h \text{ Mpc}^{-1} \lesssim k \lesssim 1 h \text{ Mpc}^{-1}$. This figure is adapted from refs. [21, 22] using https://github.com/marius311/mpk_compilation/.

spectra of tens of millions of Emission Line Galaxies (ELGs), Luminous Red Galaxies (LRGs), and quasars over the next ~ 5 years. Since ELGs comprise the largest sample of galaxies in the DESI catalogue, providing a significant fraction of the constraining power for most cosmological parameters, we only include the ELG sample in our forecasts. ELGs are active star-forming galaxies which exhibit strong nebular emission lines that originate from gas in the ionized regions around massive stars. This nebular emission includes the [OII] doublet, a distinct spectral signature that can be used to accurately determine a galaxy’s redshift. DESI’s spectrographs cover a spectral range of 360 nm to 980 nm, enabling [OII] measurements out to $z \sim 1.7$. We convert the mean values of dN_{ELG}/dz (per square degree) in Table 2.3 of ref. [23] to comoving number densities using $n_{\text{ELG}} = (180/\pi)^2 \chi^{-2} (d\chi/dz)^{-1} dN_{\text{ELG}}/dz$. These values assume a [OII] flux cut of approximately $8 \times 10^{-17} \text{ erg cm}^{-2} \text{ s}^{-1}$, and account for both selection cuts and target success rates. We follow ref. [23] and approximate the ELG linear bias as $b_{\text{ELG}} = 0.84/D(z)$, where $D(z)$ is the linear growth factor normalized to unity at $z = 0$.

The Euclid satellite [19] is a near-infrared space telescope that is currently under development, with a nominal launch date in 2022. Over its six-year long mission Euclid will measure spectra of tens of millions of galaxies out to $z \sim 2$, mapping out roughly 15,000 square degrees. Euclid consists of a 1.2 m mirror coupled to three imaging and spectroscopic instruments (with wavelength ranges 500-800 nm, 920-1250 nm, and 1250-1850 nm), enabling observations of $\text{H}\alpha$ -emitters from $0.9 < z < 1.8$. For our forecasts we use the reference case number densities $n_{\text{H}\alpha}$ and linear biases $b_{\text{H}\alpha}$ listed in Table 3 of ref. [24](see also [25]).

The Spectro-Photometer for the History of the Universe, Epoch of Reionization and Ices Explorer (SPHEREx) is a space-based, low resolution spectral survey, with a current launch date in 2024 [26]. It will image ~ 450 million galaxies in both the optical and near-infrared, primarily targeting low redshift ($z < 1$) objects and searching for evidence of primordial non-Gaussianity. We take the number densities and biases from ref. [26], which were split into five samples according to their maximum redshift uncertainty. While the redshift uncertainty can be relatively large for a subset of the combined sample, the key science goal for SPHEREx (i.e. primordial non-Gaussianity) is relatively insensitive to photo- z uncertainties. Moreover SPHEREx was designed to have a very low fraction of redshift outliers, and therefore is especially well-suited to this kind of measurement. Thus we simply add the number densities of the individual samples to get the number density of the combined sample. For the bias of the combined sample we take the appropriate weighted average: $\sum_i b_i n_i / \sum_i n_i$, where b_i and n_i are the bias and number density of the i 'th sample.

2.2 Future galaxy surveys

While the construction and deployment of various low redshift ($z < 2$) facilities are currently underway, designs for new high-redshift surveys are actively being developed. In this work we consider two of these futuristic surveys in detail: MegaMapper and MSE. These surveys would exploit dropout selection to measure galaxies out to $z \sim 5$, providing cosmological constraints that are competitive with upcoming CMB surveys (see Fig. 2).

MegaMapper [27] is a proposed highly-multiplexed spectroscopic instrument that is designed to survey galaxies at high redshift ($2 < z < 5$), covering 14,000 square degrees. Its 6.5 m telescope aperture, ~ 20 K fibers, and five year observation time would yield galaxy number densities $> 10^{-4} h^3 \text{Mpc}^{-3}$ across its redshift range. MegaMapper's primary target would be Lyman Break Galaxies (LBGs), which are abundant, well studied and well understood, representing massive, actively star-forming galaxies with a luminosity approximately proportional to their stellar mass. A sizeable fraction have bright emission lines which could facilitate obtaining a redshift [28].

For the fiducial MegaMapper experiment we use the LBG number densities listed in Table 2 of ref. [9], which are shown in Fig. 4. In §4 we also consider an idealized MegaMapper-like experiment that detects some redshift-independent fraction of LBGs \bar{n}/n_{LBG} with apparent magnitude $m < 24.5$. This is slightly unrealistic, but provides an easy way to see in what manner our forecasts are sensitive to noise levels. The ‘‘idealized’’ number densities are calculated using a Schechter function of absolute magnitude:

$$\Phi(M_{\text{UV}}) = \left(\frac{\ln 10}{2.5}\right) \phi^* 10^{-0.4(1+\alpha)(M_{\text{UV}} - M_{\text{UV}}^*)} \exp\left(-10^{-0.4(M_{\text{UV}} - M_{\text{UV}}^*)}\right). \quad (2.1)$$

We use the best fit values to the UV luminosity function for α , M_{UV}^* , and ϕ^* , which are listed in Table 3 of ref. [28]. The number density is obtained via $n_{\text{LBG}} = \int_{-\infty}^{M_c} \Phi(M) dM$, where the cutoff M_c corresponds to galaxies with apparent magnitude $m = 24.5$:

$$M_c \simeq 24.5 - 5 \log_{10} \left(\frac{D_L(z)}{10 \text{ pc}}\right) + 2.5 \log_{10}(1+z). \quad (2.2)$$

For both the fiducial and idealized MegaMapper surveys, we adopt the bias model of ref. [28], in which the bias of galaxies with apparent magnitude m is approximated to linear order in m and to second order in the scale factor:

$$b(m, z) = A(m)(1+z) + B(m)(1+z)^2 \quad (2.3)$$

with $A(m) = -0.98(m - 25) + 0.11$ and $B(m) = 0.12(m - 25) + 0.17$. We assume the LBG bias is dominated by the faintest (and most numerous) galaxies, approximating $b_{\text{LBG}}(z) = b(m = 24.5, z)$. These approximations to the number density and bias are consistent with the idealised version of MegaMapper considered in ref. [9]. As shown in Fig. 4, the fiducial MegaMapper number density is within the range $0.3 \lesssim \bar{n}/n_{\text{LBG}} < 1$, depending upon the redshift. In §4 we show that taking $\bar{n}/n_{\text{LBG}} = 0.5$ yields constraints similar to those from the fiducial MegaMapper survey.




Figure 4. *Left:* In blue we plot several values of \bar{n}/n_{LBG} , where n_{LBG} is the idealized LBG number density with $m < 24.5$. For comparison, we plot the number density of the fiducial MegaMapper survey in orange. *Right:* Comoving volume per dz per sky area ($\chi^2 dX/dz$) as a function of redshift. The shaded blue, purple, and (hatched) grey regions represent the volume covered by Euclid, MegaMapper and PUMA respectively, excluding sky coverage.

The Maunakea Spectroscopic Explorer (MSE) [29], planned for first light in 2029, will cover over 10,000 square degrees. MSE will couple a 11.25 m mirror with a 1.5 square degree field of view to ~ 4000 fibers, feeding to spectrographs that cover 360 to 1300 nm. This design enables the detection of ELGs out to $z \sim 2.4$, and LBGs between $2.4 < z < 4$. For the ELG sample we use the number densities listed in Table 1 of ref. [30], and take $b_{\text{ELG}}(z) = D(0)/D(z)$, where $D(z)$ is the linear growth factor. To be consistent with our assumptions regarding LBGs, we use the same linear bias $b(m = 24.5, z)$ for both MegaMapper and the MSE LBG sample. For the LBG sample we estimate the number density via $\int_{-\infty}^{M_c} \Phi(M)E(M)dM$, where Φ and M_c are defined in Eqs. (2.1) and (2.2) respectively, and $E(M)$ is MSE’s LBG target efficiency. We take the efficiency to be an average over the templates shown in Figure 2 of ref. [30], assuming that 40% of the LBGs have a negative equivalent width (EW), 30% have $0 < \text{EW} < 20$ and 30% have $\text{EW} > 20$.

SpecTel [31, 32] is a proposed spectroscopic survey in the southern hemisphere, making it ideally situated for cross-correlations with SO or CMB-S4. SpecTel would couple an 11.4 m dish (with a 5 square degree FoV) to 60,000 fibers – enabling 120 million fiber exposures (each $\geq 4,000$ seconds long) over its survey period – which feed to spectrographs covering 360 to 1330 nm. Its design would permit observations of LBGs and Lyman Alpha Emitters (LAEs) out to redshift ~ 5 , with number densities 2 to 5 times those achievable by a MegaMapper-like survey, depending upon the redshift bin. Since the fiducial survey design is still in flux, we do not perform separate forecast for SpecTel. However, in §4 we explore how MegaMapper’s parameter constraints improve with increased LBG number density, which can be interpreted as a rough estimate for SpecTel constraints.

The Roman telescope [33] is a space-based NASA mission set to launch in 2025, carrying a spectrometer that covers the visible to near-infrared spectrum. We have not attempted

forecasts for Roman because the survey design is still not finalized. Under the straw-man proposal for a $2,000 \text{ deg}^2$ survey, constraints from Roman would not be competitive with those from DESI or the other proposed surveys we investigate in detail. Study of trade-offs which would improve Roman’s performance in these forecasts is beyond the scope of this paper and would require a better understanding of achievable depth and area for the galaxy redshift survey given its current constraints.

2.3 21-cm intensity mapping

Proof-of-concept interferometers dedicated to 21-cm mapping are currently being deployed [34–36], while designs for future larger-scale facilities are actively under development [13, 37, 38]. We consider two 21-cm surveys in detail: the Hydrogen Intensity and Real-time Analysis eXperiment (HIRAX) [37] and the Packed Ultra-Wideband Mapping array (PUMA) [13]. HIRAX is a 400–800 MHz radio interferometer that is currently under development for deployment in South Africa, while PUMA is a proposed ultra-wideband, low-resolution and transit interferometric radio telescope operating at 200–1100 MHz. PUMA’s fiducial telescope design is a hexagonal array of 32K 6 m dishes, designed to survey half of the sky ($f_{\text{sky}} = 0.5$) over five years. We also consider a smaller, proof-of-concept design with only 5K dishes (PUMA-5K). In our noise model (see below), we assume a hex-packed array with a 50% fill factor. This fill factor is accounted for by doubling the number of detectors and setting the observation time to $t_{\text{int}} = 5/4$ years. All other couplings, efficiencies, temperatures etc are set to the values quoted in ref. [39]. We use the same noise model for HIRAX, but with a 1K, fully filled array, and $t_{\text{int}} = 4$ years.

Below $z \simeq 6$ the majority of the neutral Hydrogen in the Universe resides in dense, self-shielded regions within galaxies [40]. While the details of how HI traces the dark matter are not well understood, current observations and numerical simulations suggest that we can approximate the HI mass within a given halo with $M_{\text{HI}}(M_h; z) = A(z)M_h e^{-M_{\text{min}}/M_h}$, where $M_{\text{min}} = 5 \times 10^9 h^{-1} M_{\odot}$ [41]. Using this relation, along with a halo number density and bias function, we can then approximate the linear bias of HI by weighting the halo bias with the halo energy density:

$$b_{\text{HI}}(z) = \frac{\int_0^{\infty} b(M_h; z) n(M_h; z) M_{\text{HI}}(M_h; z) dM_h}{\int_0^{\infty} n(M_h; z) M_{\text{HI}}(M_h; z) dM_h}. \quad (2.4)$$

We adopt the same conventions as ref. [41] and use the Sheth-Tormen [42] halo bias and mass function.

There are two sources of noise in HI observations: a shot noise due to the finite sampling of the continuous HI field, and the thermal noise from the instrument. For all HI surveys we adopt the thermal noise model described in the appendix of ref. [39] (see also [43, 44]), along with the shot noise from ref. [45, 46]. Further discussion of observational and data-analysis issues related to interferometric, 21-cm intensity mapping can be found in ref. [47].

2.4 CMB experiments

In addition to LSS probes, we consider external priors from CMB surveys, with configurations similar to the Planck [21], Simons Observatory (SO; [48]) CMB-S4 [14] and LiteBIRD [49] surveys. To obtain the Planck Fisher matrix, we use the bands, beams and sensitivities from Table 4 of ref. [21]. We include both temperature and polarization anisotropies, and we adjust the minimum $\ell_{\text{min,P}}$ in polarization to match the published value of uncertainty

on the optical depth $\sigma(\tau) = 0.007$. We have checked that this method gives good agreement with the published uncertainties on the standard cosmological parameters, while noting that the main purpose of this paper is not forecasting parameters for the next generation of CMB experiments, but rather to highlight the improvements from high-redshift LSS surveys. Because of this, we consider a simplified setup in which detector noise is treated as white noise uncorrelated between different frequency channels, and we do not include any effects from the atmosphere on large scales (beyond the scale cut $\ell_{\min} = 30$), and do not include any instrumental or foreground nuisance parameters. For SO and CMB-S4 we take $\ell_{\max,T} = 3000$ in temperature and $\ell_{\max,P} = 5000$ in polarization, and $\ell_{\min,T,P} = 30$ for both. Since SO and CMB-S4 are ground based and not expected to measure the very largest scales due to contamination from the atmosphere, we always add a low- ℓ Planck (or LiteBIRD) Fisher matrix to SO or CMB-S4. This is accomplished by splitting the SO/S4/Planck forecasts into high- and low- ℓ Fisher matrices (with cutoff $\ell = 30$), and adding the low- ℓ Planck Fisher matrix to the high- ℓ SO/S4 Fisher matrix. This simplistic addition is justified since the two Fisher matrices use a disjoint set of multipoles, and are therefore largely uncorrelated (we expect any mode coupling from lensing or non-Gaussian foregrounds to be small). A fiducial model for the foregrounds is added to the total power spectrum in temperature, following the model of ref. [50] (with the modifications listed in ref. [48]). We neglect any foregrounds in polarization, noting that point source masking should be very effective at mitigating their effect on small scales. Unless stated otherwise, we use the 2-point function of the unlensed primary CMB temperature and polarization fluctuations for the derivatives with respect to the cosmological parameters, and do not include CMB lensing reconstruction. This allows us to simply add the Fisher matrices for CMB and LSS probes. Another feature of using the unlensed CMB power spectrum is that it should provide accurate forecasts from high- ℓ temperature. If we included lensing the damping tail would be proportional to the lensing amplitude, and therefore largely degenerate with the foreground amplitude and shape, as well as baryon effects in the lensing power spectrum which is not marginalized over here. This can lead to artificially tight constraints [51]. We also consider a futuristic LiteBIRD-like space based CMB experiment, capable of reaching the cosmic variance limit $\sigma(\tau) \approx 0.002$ from $\ell < 30$. For this we take a $1.4\mu\text{K-arcmin}$ experiment in polarization with a 33 arcmin Gaussian beam covering 65% of the sky. Although not considered in this work, we note that an instrument like PICO [52] would also be highly complementary to LSS probes.

Experiment	z_{\min}, z_{\max}	Sample	f_{sky}
DESI [23]	0.6, 1.7	ELG	0.34
Euclid [19, 24, 53]	0.9, 1.8	H α	0.36
SPHEREx [26, 54]	0.1, 4.3	low- z galaxies	0.65
MegaMapper [27]	2, 5	LBG	0.34
MSE [29]	1.6, 4	ELG + LBG	0.24

Experiment	z_{\min}, z_{\max}	Sample	f_{sky}
PUMA [13]	2, 6	21-cm	0.50
HIRAX [37]	0.8, 2.5	21-cm	0.36

Table 1. *Left:* A list of the spectroscopic galaxy surveys considered in our forecasts, along with their redshift ranges, samples and sky coverage. *Right:* A list of 21-cm experiments, their redshift range and sky coverage.

3 Forecasting formalism

In this section we outline our forecasting procedure, based on the Fisher formalism. We examine three primary observables of interest: primary anisotropies in the CMB (as described

in §2.4), the redshift space galaxy power spectrum (§3.1), and CMB lensing (§3.2). We take $\Lambda\text{CDM}+M_\nu$ as our fiducial cosmology, with ΛCDM values given in Table 2 and $N_{\text{eff}} = 3.046$, consistent with refs. [21, 55]. We assume a normal neutrino hierarchy with minimal mass: $m_{\nu_1} = 0$, $m_{\nu_2} = 0.01$ eV and $m_{\nu_3} = 0.05$ eV. Unless stated otherwise the neutrino mass $M_\nu = \sum_i m_{\nu_i}$ and N_{eff} are held fixed in all forecasts. A list of all of the free parameters in our base model is given in Table 2.

3.1 Full shape measurements

Both historically and currently, most forecasts use the linear power spectrum and scale-independent bias, or the non-linear matter power spectrum computed using fitting formulae, to model the power spectrum of a biased tracer [13, 23, 24, 46, 56–62]. In this work, we use more sophisticated models based on perturbation theory and a general bias expansion (see refs. [20, 63–65] for related work on forecasts and ref. [66] for inclusion of scale-dependent biases in the halo model). We consider a general quadratic bias model and non-linear evolution up to one-loop order in perturbation theory. On large scales and at high redshifts, perturbation theory provides an accurate model that allows us to take into account scale-dependent bias, broadening of the acoustic peaks and mode-coupling due to non-linear evolution. Including these effects should improve the reliability of the forecasts. We further discuss the differences in linear vs. non-linear forecasting in §4.11.

To include the non-linear effects and more complex bias model we use the `velocileptors` LPT_RSD module [8, 41, 67], which takes the linear power spectrum and biases as inputs, and self-consistently calculates the non-linear redshift-space power spectrum $P_{gg}^{1\text{-loop}}(k, \mu)$ to one-loop order in perturbation theory. We use the linear CDM+baryon power spectrum $P_{cb}^{\text{lin}}(k, z)$ as our input to `velocileptors`, calculated using CLASS [68], along with a second order biasing scheme:

$$\delta_g(\mathbf{x}, z) = b \delta_{cb}(\mathbf{x}, z) + \frac{1}{2} b_2 \delta_{cb}^2(\mathbf{x}, z) + b_s s^2(\mathbf{x}, z) \quad (3.1)$$

where δ_g is the non-linear tracer density contrast, δ_{cb} is the non-linear CDM+baryon density contrast, and $s_{ij} = (\partial_i \partial_j / \partial^2 - \delta_{ij} / 3) \delta_{cb}$ is the shear field. To better connect our formalism to previous literature here, and elsewhere in the text, these biases are in the Eulerian frame³. However, the LPT_RSD module computes $P_{gg}^{1\text{-loop}}(\mathbf{k})$ using Lagrangian perturbation theory, which has been shown to perform well in these situations [67] and allows for a simple treatment of models with complex features in the power spectrum [8]. To compute the Lagrangian biases, which are inputs to `velocileptors`, we perform the following rotation [69, 70]:

$$b^L = b - 1 \quad b_2^L = b_2 - \frac{8}{21}(b - 1) \quad b_s^L = b_s + \frac{2}{7}(b - 1). \quad (3.2)$$

We use the linear biases given in §2. The fiducial values of b_2 and b_s are set by assuming $b_2^L = b_s^L = 0$. This choice is only a rough approximation whose validity is pending more data. In all forecasts we marginalize over these non-linear biases, making our results mostly insensitive to their fiducial values (see Appendix E for a more detailed discussion).

To model the effects of small-scale physics, we include a handful of counterterms (α_n) and stochastic contributions (N_n) in our fiducial power spectra:

$$P_{gg}(\mathbf{k}) = P_{gg}^{1\text{-loop}}(\mathbf{k}) + (\alpha_0 + \alpha_2 \mu^2 + \alpha_4 \mu^4) (k/k_*)^2 P_{cb}^{\text{Zel}}(\mathbf{k}) + N_0 + N_2(\mu k)^2 + N_4(\mu k)^4, \quad (3.3)$$

³We assume cubic bias operators are fixed and generated by time evolution in Lagrangian perturbation theory.

where N_2 and N_4 encode small-scale velocities (i.e. FoG effects and redshift errors, if appropriate), N_0 is the shot noise, $k_* = 1 h \text{Mpc}^{-1}$, P_{cb}^{Zel} is the CDM+baryon power spectrum in the Zel'dovich approximation, and $\mu = \hat{\mathbf{k}} \cdot \hat{\mathbf{n}}$ is the cosine of the angle between the wavevector and the line of sight. Throughout we approximate $P_{gg}(k, \mu) \simeq \sum_{\ell=0}^4 \mathcal{L}_\ell(\mu) P_{gg,\ell}(k)$ by its first three non-vanishing multiples, where $\mathcal{L}_\ell(\mu)$ is the Legendre polynomial of degree ℓ . For the surveys, redshifts and k -ranges considered in this work we find that this approximation is good to 2%.

We determine the fiducial value of α_0 by fitting the `velocileptors` output to CLASS' default version of `HaloFit` [71]. The fiducial values for the higher order stochastic terms (α_2, α_4) are set to zero purely for convenience. We take the Poisson value $\bar{n}_{\text{gal}}^{-1}$ for the fiducial value of N_0 and set $N_2 = -N_0 \sigma_v^2$, where σ_v is a typical velocity dispersion. For galaxy surveys we set the velocity dispersion to 100 km/s, converted to comoving length units: $\sigma_v = (1+z)(100 \text{ km/s})/H(z)$. Finger-of-god effects in 21-cm maps are much smaller, driven by a small number of satellites (occupying $\sim 10\%$ of the sample) with typical rms velocities $\approx 40 \text{ km/s}$ [72]. Thus for 21-cm surveys, we choose a fiducial velocity dispersion of 10 km/s ($\sim \sqrt{0.1} 40 \text{ km/s}$). The fiducial value of N_4 is set to zero purely for convenience. In all forecasts we marginalize over these counterterms and stochastic contributions, making our results largely insensitive to their fiducial values, as discussed in Appendix E.

The modeling of the power spectrum for 21-cm surveys is slightly more complex than for galaxy surveys. Radio interferometers measure the 21-cm signal in intensity units, so that the power spectrum has units of $(\text{temperature})^2 \times \text{volume}$:

$$P_{\text{HI}}^{\text{obs}}(\mathbf{k}) = T_b^2 [P_{\text{HI}}(\mathbf{k}) + P_{\text{SN}}] + P_{\text{N}}, \quad (3.4)$$

where $P_{\text{HI}}(\mathbf{k})$ is the non-linear HI power spectrum (taking the same form as Eq. 3.3, with P_{SN} replacing N_0), P_{SN} is the shot-noise coming from halos that contribute to the HI signal, P_{N} is the instrumental noise [39], and T_b is the mean 21-cm brightness temperature [73]:

$$T_b(z) = 188(1+z)^2 E(z) \Omega_{\text{HI}}(z) \text{ mK}, \quad (3.5)$$

where $E(z) \equiv H(z)/H_0$ and $\Omega_{\text{HI}}(z) \equiv \rho_{\text{HI}}(z)/\rho_c(z)$ is the cosmic density of neutral hydrogen, taken from the fitting formula of ref. [74]. In all forecasts that include 21-cm surveys we marginalize over the 21-cm brightness temperature by default as its true value is still largely uncertain.

Because our fiducial model includes massive neutrinos, we need to specify their perturbative treatment in our setup. It is by now well established that halos and galaxies can be well approximated as biased tracers of the dark matter and baryon fluids only [75–81]. For this reason our real-space power spectrum calculation goes through as for the no-neutrino case if we make the trivial replacement of the total matter field with P_{cb} , the dark matter plus baryon spectrum, as described above. We also assume the Green's function associated with higher order perturbative solutions is very well approximated by the Einstein-de-Sitter one, which implies the structure of the loops remains unchanged. This is an excellent approximation for the small fiducial value of neutrino masses considered in this work [82]. In principle one needs to consider the scale-dependence of the growth rate, which we shall neglect for the same reason mentioned above.

For our full shape forecasts we work at the power spectrum level. As in Eq. (3.3), let $P_{gg}(k, \mu, z)$ denote the observed nonlinear redshift-space power spectrum of the matter tracer. Let i index a (μ_i, k_i) pair, e.g. $P_{gg,i}(z) = P_{gg}(k_i, \mu_i, z)$. We divide each survey into linearly

spaced redshift bins $\{B_m; m = 1, 2, \dots\}$, and calculate the corresponding Fisher matrix in each bin as⁴:

$$F_{ab}(B_m) = \sum_{ij} \frac{\partial P_{gg,i}(z)}{\partial \theta_a} C_{ij}^{-1}(z) \frac{\partial P_{gg,j}(z)}{\partial \theta_b} \Big|_{z=z_m} \quad (3.6)$$

$$\text{where } C_{ij}(z) = \delta_{ij}^K \frac{4\pi^2}{k_i V_m d\mu d \ln k} [P_{gg,i}(z)]^2.$$

Here V_m is the comoving volume of bin B_m , which implicitly includes a factor of f_{sky} , $z_m = (z_{\text{max}} + z_{\text{min}})/2$ is the central redshift in bin B_m , and θ_a represents either a cosmological or a nuisance parameter. Throughout we assume that $P_{gg}(k, \mu, z)$ is even in μ , and use Simpson's rule to evaluate Eq. (3.6) with 2,000 logarithmically spaced k -bins (ranging from 5×10^{-4} to $1 h \text{Mpc}^{-1}$) and 100 linearly spaced μ -bins (ranging from 0 to 1). We limit the domain of integration by multiplying the summand by an appropriate window function. To calculate the Fisher matrix of the full survey we make the standard approximation that correlations between non-overlapping redshift bins are negligible, so that the final Fisher matrix is obtained by summing over redshift bins:

$$F_{ab} = \sum_m F_{ab}(B_m). \quad (3.7)$$

In Eq. (3.6) we neglect any temporal or spatial variation in the number density within each bin. This effect can be included analytically [83], and can lead to significant error in the forecast for sufficiently large redshift bins, or for sufficiently low number densities. However, as we show in Appendix B, for the surveys and redshift bins considered in this work these effects are negligibly small. We thus ignore this correction, and take Eq. (3.6) at face value.




Figure 5. *Left:* $k_{\text{max}}(z)$ and $\ell_{\text{max}}(z) = k_{\text{max}}(z)\chi(z)$, assuming that k_{max} is some constant fraction of $k_{\text{nl}}(z) = \Sigma^{-1}(z)$. *Right:* Domain of integration in the $k_{\parallel} - k_{\perp}$ plane at $z \simeq 3.9$. The grey region is excluded by the $k < k_{\text{nl}}$ constraint. For 21-cm surveys the blue region is excluded by the (pessimistic) foreground wedge. The red region is excluded by FoG-like effects, chosen so that $N_0 \sigma_v^2(k\mu)^2$ is less than 20% of the fiducial power. In this figure we plot the red region for a MegaMapper-like survey (it would shift to higher k_{\parallel} for a 21-cm survey).

In all forecasts we take $k_{\text{min}} = \max[0.003 h \text{Mpc}^{-1}, 2\pi/V_m^{1/3}]$ for each redshift bin. Since the clustering signal is low for $k < k_{\text{min}}$, our results would not change significantly if a

⁴The shot noise $N_0(z) = 1/n_{\text{gal}}(z)$ is included in our definition of the observed power spectrum, so it does not appear explicitly in our covariance.

slightly different k_{\min} were assumed, with the exception of local primordial non-Gaussianity, as discussed in Sec. 4.6. Unless explicitly stated otherwise we take $k_{\max}(z) = k_{\text{nl}}(z)$, where $k_{\text{nl}}(z) = \Sigma^{-1}(z)$ is the non-linear scale set by the RMS displacement in the Zel’dovich approximation:

$$\Sigma^2(z) = \int \frac{dk}{6\pi^2} P^{\text{lin}}(k, z), \quad (3.8)$$

where $P^{\text{lin}}(k, z)$ is the linear matter power spectrum. For 21-cm surveys, we also include a foreground-wedge [39] that further constrains the limits of integration:

$$k_{\parallel} > \max \left[k_{\perp} \frac{\chi(z)H(z)}{c(1+z)} \sin(\theta_w(z)), k_{\parallel}^{\min} \right] \quad \text{where} \quad \theta_w(z) = N_w \frac{1.22}{2\sqrt{0.7}} \frac{\lambda_{\text{obs}}}{D_{\text{phys}}} \quad (3.9)$$

where $\lambda_{\text{obs}} = (1+z)21$ cm is the observed HI wavelength, D_{phys} is the physical diameter of the dish, $k_{\parallel} = k\mu$ and $k_{\perp} = k\sqrt{1-\mu^2}$. We take $D_{\text{phys}} = 6$ m for both HIRAX and PUMA. For each 21-cm survey we consider ‘‘optimistic’’ and ‘‘pessimistic’’ foreground cases, in which $N_w = 1$, $k_{\parallel}^{\min} = 0.01 h \text{ Mpc}^{-1}$ and $N_w = 3$, $k_{\parallel}^{\min} = 0.1 h \text{ Mpc}^{-1}$ respectively. This is consistent with the definitions in ref. [39]. Finally, we note that non-perturbative effects like Fingers-of-God (FoG) limit the information that we can extract from high values of μ . The importance of FoG is not known a priori, but roughly speaking they can be accounted for by a stochastic term $N_0\sigma_v^2(k\mu)^2$ (see Eq. 3.3) on sufficiently large scales. We therefore make the conservative choice of integrating only over modes where this term is less than 20% of the fiducial power $P_{gg}(k, \mu)$. We have checked that this choice in cutoff has an insignificant impact on our forecasts, as discussed in Appendix E and shown in Fig. 28. In Fig. 5 we show both the foreground wedge (for the pessimistic case) and the region excluded by FoG-like effects for our fiducial MegaMapper survey.

3.2 CMB lensing and cross-correlations

Weak lensing of the CMB is a tracer of the matter distribution integrated along the line of sight [84], and is one of the central science cases for upcoming CMB surveys [48, 49]. Throughout we work under both the Born and Limber approximations so that C_{ℓ}^{XY} may be expressed as a single integral:

$$C_{\ell}^{XY} = \int_0^{\chi_*} d\chi \frac{W^X(\chi)W^Y(\chi)}{\chi^2} P_{\delta_X\delta_Y} \left(k_{\perp} = \frac{\ell + 1/2}{\chi}, k_{\parallel} = 0 \right), \quad (3.10)$$

where X and Y denote either the CMB lensing convergence κ or the galaxy sample g , and $\delta_{\kappa} = \delta_m$. In Eq. (3.10) we have included the lowest order correction to the Limber approximation, replacing $\ell \rightarrow \ell + 1/2$, which increases the accuracy to $\mathcal{O}(\ell^{-2})$ [85]. In principle both approximations could be relaxed, but in practice we do not expect them to have a large impact on the parameters constraints [86–92]. The projection kernels take the form:

$$W^{\kappa}(\chi) = \frac{3}{2}\Omega_m H_0^2(1+z) \frac{\chi(\chi_* - \chi)}{\chi_*} \quad \text{and} \quad W^g(\chi) \propto H(z) \frac{dN}{dz}, \quad (3.11)$$

where W^g is normalized so that $\int d\chi W^g(\chi) = 1$, and χ_* is the comoving distance to the surface of last scattering. For simplicity we neglect any contribution from magnification bias, as this will be measured by future LSS surveys to high enough accuracy to be treated as an essentially fixed parameter [41].

The nonlinear real space power spectra (P_{gg} and P_{gm}) are calculated to one-loop in LPT using `velocileptors'` `cleft_fftw.py`, which includes a second order biasing scheme (b, b_2, b_s), two counterterms (α_0 and α_x) for the auto- and cross-correlations, and a shot noise contribution to P_{gg} . Note that lensing cross-correlations introduces only one new nuisance term (α_x), and that P_{gm} and P_{gg} manifestly have different dependencies on the bias parameters. Thus lensing cross-correlations play a role in breaking degeneracies, improving the constraints on parameters which are highly degenerate with bias (e.g. M_ν), as discussed in Appendix D (see Fig. 27). The fiducial values for the biases, counterterms and shot noise are identical to those for the full shape power spectrum, and are listed in Table 2.

For both the cross-correlation $C_\ell^{\kappa g}$ and the galaxy auto-correlation C_ℓ^{gg} we approximate Eq. (3.10) with 100 integration points that are linearly spaced in the redshift interval corresponding to the galaxy sample g . In principle the power spectra should be evaluated at each integration point, but this is computationally expensive to do in practice. To speed up our calculations we split the redshift interval into m linearly-spaced bins, each of which contain $100/m$ integration points. Within each bin the power spectrum is evaluated at a single effective redshift [93]:

$$z_{\text{eff}}^{XY} = \frac{\int d\chi W^X(\chi)W^Y(\chi)z(\chi)/\chi^2}{\int d\chi W^X(\chi)W^Y(\chi)/\chi^2}, \quad (3.12)$$

which is chosen to cancel the linear order correction in the expansion of $P_{\delta_X\delta_Y}(k, z)$ about z_{eff}^{XY} . We choose m to be the smallest integer that yields bins with widths $\Delta z \leq 0.2$, resulting in $\mathcal{O}(0.1\%)$ accuracy for $30 < \ell < 500$ and $1 < z < 6$. Achieving a similar accuracy for the convergence power spectrum via direct integration is computationally inefficient. We instead calculate $C_\ell^{\kappa\kappa}$ with CLASS, using its default version of HALOFIT to model nonlinearities.

We consider both SO and S4-like experiments when providing lensing constraints. For SO we use the noise N_ℓ^κ from the [SO noise calculator](#) while for S4, we use the lensing noise curve described on the [CMB-S4 Wiki](#). These curves are plotted in Fig. 6. Both of these noises include atmospheric contributions, and are calculated with an iterated minimum variance (MV) quadratic estimator and MV Internal Linear Combination (ILC) of both CMB temperature ($30 < \ell_T < 3000$) and polarization ($30 < \ell_P < 5000$) data.

In our forecasts we work at the power spectrum level and split our galaxy sample g into n non-overlapping redshift bins, denoted by g_1, \dots, g_n , so that our averaged data vector is $\boldsymbol{\mu}_\ell^T = (C_\ell^{\kappa\kappa}, C_\ell^{\kappa g_1}, \dots, C_\ell^{\kappa g_n}, C_\ell^{g_1 g_1}, \dots, C_\ell^{g_n g_n})$. The corresponding Fisher matrix is

$$F_{ij} \simeq \frac{f_{\text{sky}}}{2} \sum_{\ell=\ell_{\text{min}}}^{\ell_{\text{max}}} (2\ell + 1) \frac{\partial \boldsymbol{\mu}_\ell^T}{\partial \theta_i} \mathbf{C}_\ell^{-1} \frac{\partial \boldsymbol{\mu}_\ell}{\partial \theta_j}, \quad (3.13)$$

where the elements of the covariance \mathbf{C}_ℓ are given by⁵

$$(\mathbf{C}_\ell)_{XY} = \begin{cases} (C_\ell^{\kappa\kappa} + N_\ell^\kappa)^2 & X = Y = \kappa\kappa \\ \frac{1}{2} \left[(C_\ell^{\kappa\kappa} + N_\ell^\kappa) C_\ell^{g_i g_i} \delta_{ij}^K + C_\ell^{\kappa g_i} C_\ell^{\kappa g_j} \right] & X = \kappa g_i, Y = \kappa g_j \\ (C_\ell^{g_i g_i})^2 \delta_{ij}^K & X = g_i g_i, Y = g_j g_j \\ (C_\ell^{\kappa\kappa} + N_\ell^\kappa) C_\ell^{\kappa g_i} & X = \kappa\kappa, Y = \kappa g_i \\ (C_\ell^{\kappa g_i})^2 & X = \kappa\kappa, Y = g_i g_i \\ C_\ell^{g_i g_i} C_\ell^{\kappa g_i} \delta_{ij}^K & X = g_i g_i, Y = \kappa g_j \end{cases} \quad (3.14)$$

⁵As in §3.1, we include the shot noise in our fiducial $C_\ell^{g_i g_i}$, so it does not explicitly appear in our covariance.

In all forecasts we take $\ell_{\min} = 30$. To remove information from modes that are too small to be accurately modeled using PT, we multiply the appropriate elements of the covariance matrix ($C_\ell^{g_i g_i}$ and $C_\ell^{\kappa g_i}$) by 10^{20} whenever $\ell > k_{\text{nl}}(z_i)\chi(z_i)$, where z_i is central redshift in the i 'th redshift bin. Throughout we assume full spatial overlap of CMB and LSS surveys for cross-correlations: $f_{\text{sky}}^{\text{cross}} = \min(f_{\text{sky}}^{\text{LSS}}, f_{\text{sky}}^{\text{CMB}})$, with $f_{\text{sky}}^{\text{CMB}} = 0.4$ for a SO-like experiment.




Figure 6. *Left:* Lensing noise N_ℓ^κ for Planck, Simons Observatory and CMB-S4. The convergence power spectrum is shown in black. *Middle:* The galaxy auto-correlation $C_\ell^{g_i g_i}$ for MegaMapper, where g_i is the LBG sample in the i 'th redshift bin: $z \in [2.0, 2.8], [2.8, 3.5], [3.5, 4.3],$ or $[4.3, 5.0]$. *Right:* The convergence-galaxy cross-spectra in each redshift bin.

Due to projection effects, the lensing auto-power spectrum receives substantial contributions from small scales, beyond the reach of perturbation theory. For such modes, baryonic effects cannot be neglected either, further complicating the model predictions. To be conservative we impose a cut ($\ell < 500$) on the lensing angular modes included in our forecasts, but note that this is active area of research and improvements are likely to be possible by the time the data arrive [94–97]. For $\ell < 500$, baryonic feedback contributes at the subpercent level to the convergence power spectrum [98], while at $\ell = 500$, only $\sim 15\%$ of the power comes from $k > 0.2 h \text{ Mpc}^{-1}$.

3.3 Combined forecasts

In our forecasts we treat the CMB surveys described in §2.4 as priors on ΛCDM , M_ν , N_{eff} and Ω_k . Since our CMB priors are calculated using *unlensed* power spectra, we neglect any covariances between these priors and lensing or full shape data. Thus to combine a CMB prior with a LSS Fisher matrix \mathbf{F}_{LSS} with basis $\{\Lambda\text{CDM}, M_\nu, N_{\text{eff}}, \Omega_k, \dots\}$, we simply add the CMB Fisher matrix to the appropriate submatrix of \mathbf{F}_{LSS} .

When combining lensing⁶ ($C_\ell^{\kappa\kappa}, C_\ell^{\kappa g_i}, C_\ell^{g_i g_i}$) and full shape clustering data ($P_{gg}(\mathbf{k})$) from the same survey, we impose a $k_{\parallel} > 10^{-3} h \text{ Mpc}^{-1}$ cut on the full shape Fisher matrix. Since lensing probes modes with $k_{\parallel} \sim 0$, this cut nulls any covariance between the two observables in linear theory. We therefore neglect any covariance between the lensing and full shape data, and simply add the respective Fisher matrices when showing combined constraints.

We also neglect any (negligibly small) covariances between two non-overlapping LSS surveys when combining their full shape data. The same cannot be done when combining lensing data from two surveys, however, since the covariance between $C_\ell^{\kappa g_i}$ and $C_\ell^{\kappa g_j}$ is nonzero even when the tracers g_i and g_j are at different redshifts. To avoid this complication we never

⁶From here on out we abbreviate $(C_\ell^{\kappa\kappa}, C_\ell^{\kappa g_i}, C_\ell^{g_i g_i})$ as “lensing data”, “+ lensing” or simply “lensing”, unless stated otherwise.

Parameter	Definition	Fiducial value
h	Hubble parameter: $H_0/100$ km/s/Mpc	0.677
$\ln(A_s)$	$\ln(\text{primordial amplitude})$	-19.97
n_s	spectral index	0.96824
ω_c	fractional dark matter density: $\Omega_c h^2$	0.11923
ω_b	fractional baryon density: $\Omega_b h^2$	0.02247
τ	optical depth to reionization	0.0568
b	linear Eulerian bias: $\delta_g \supset b\delta_{cb}$	survey dependent
b_2	quadratic Eulerian bias: $\delta_g \supset \frac{1}{2}b_2\delta_{cb}^2$	$(8/21) \times (b - 1)$
b_s	shear bias: $\delta_g \supset b_s s^2$	$-(2/7) \times (b - 1)$
α_0	$P_{gg} \supset \alpha_0 (k/k_*)^2 P_{cb}^{\text{Zel}}$	$1.22 + 0.24 b^2 (z - 5.96)$
α_{2n}	$P_{gg} \supset \alpha_{2n} \mu^{2n} (k/k_*)^2 P_{cb}^{\text{Zel}}, n = 1, 2$	0
N_0	Poisson white noise: $P_{gg} \supset N_0$	survey dependent
N_2	FoG-like contribution: $P_{gg} \supset N_2 (k\mu)^2$	survey dependent
N_4	FoG-like contribution: $P_{gg} \supset N_4 (k\mu)^4$	0
T_b	21-cm brightness temperature	Eq. (3.5)
α_x	$P_{gm} \supset \alpha_x (k/k_*)^2 P_{cb}^{\text{Zel}}$	0

Table 2. A list of the free parameters in our base model, which are marginalized over in all forecasts unless explicitly stated otherwise. In the first set of rows we list the Λ CDM parameters and their fiducial values. In the second set of rows we list the nuisance terms that we use to model redshift-space clustering. The last row shows the counterterm introduced when including lensing cross-correlations. Here δ_g is the non-linear density contrast of the tracer with power spectrum P_{gg} , δ_{cb} is the non-linear density contrast of CDM+baryons, P_{cb}^{Zel} is the power spectrum of the CDM+baryons within the Zeldovich approximation, $s_{ij} \equiv (\partial_i \partial_j / \partial^2 - \delta_{ij}/3)\delta_{cb}$ is the shear field, $s^2 \equiv s_{ij}s_{ij}$, and $k_* = 1 h \text{ Mpc}^{-1}$.

combine lensing data from two different surveys – when combining full shape information from two spectroscopic surveys with lensing data, we only include the lensing data from the higher redshift survey. Since the high- z survey (e.g. MegaMapper) typically dominates the constraining power, neglecting the lensing-galaxy cross-correlations of the low- z survey (e.g. DESI) has a negligible impact on our results. Since the foreground wedge (§3.1) restricts access to modes with $k_{\parallel} \sim 0$, we never include 21-cm \times CMB lensing in our results. Thus when adding lensing to (low- z spectroscopic survey) + (21-cm) clustering data, we only include the lensing data from the low- z survey.

The full set of free parameters in our base model is given in Table 2. In all forecasts that include full shape information we marginalize over $\{\Lambda\text{CDM}, b, b_2, b_s, \alpha_0, \alpha_2, \alpha_4, N_0, N_2, N_4\}$, while in all forecasts that include lensing we marginalize over $\{\Lambda\text{CDM}, b, b_2, b_s, \alpha_0, N_0, \alpha_x\}$. We marginalize over the 21-cm brightness temperature T_b in any forecast that includes 21-cm full shape data. We treat all time-dependent nuisance terms (biases, counterterms, stochastic terms, T_b) in different redshift bins as separate, independent variables which are individually marginalized over. We do the same for nuisance terms of separate surveys when showing combined constraints.

3.4 Derivatives

With the exception of derivatives with respect to the three stochastic terms (N_0, N_2, N_4) and the 21-cm brightness temperature, all of our derivatives are calculated numerically. For




Figure 7. Some exemplary logarithmic derivatives at $z = 2.5$, $\mu = 0$. The solid curves are calculated using Lagrangian perturbation theory, the dashed curves are calculated using linear theory. These derivatives do not account for the A-P (§3.5) effect.

parameters whose fiducial value, θ_{fid} , is within the physically allowed region we use a five-point, two-sided derivative with relative step size $\delta\theta/\theta_{\text{fid}}$, which takes the form:

$$\left. \frac{\partial P}{\partial \theta} \right|_{\theta_{\text{fid}}} = \frac{-P_{+2} + 8P_{+1} - 8P_{-1} + P_{-2}}{12\delta\theta} + \mathcal{O}(\delta\theta^4). \quad (3.15)$$

where $P_{\pm n} \equiv P(\theta_{\text{fid}} \pm n\delta\theta)$. When the parameter has its fiducial value at a boundary of the allowed region we instead use a five-point, one-sided derivative:

$$\left. \frac{\partial P}{\partial \theta} \right|_{\theta_{\text{fid}}} = \frac{-3P_{+4} + 16P_{+3} - 36P_{+2} + 48P_{+1} + 25P_{+0}}{12\delta\theta} + \mathcal{O}(\delta\theta^4). \quad (3.16)$$

We take $\delta\theta/\theta_{\text{fid}} = 0.01$ for most parameters. When the fiducial value of the parameter is zero, we assume $\delta\theta = 0.01$ by default. τ , M_ν and A_{lin} (see §4.7) require larger (or smaller) step sizes for numerical convergence: $\delta\tau = 0.3\tau$, $\delta M_\nu = 0.05M_\nu$ and $\delta A_{\text{lin}} = 0.002$ respectively. For these choices the convergence of the derivatives is better than 0.5% for all parameters. We show some illustrative examples in Fig. 7.

When taking derivatives with respect to M_ν , we fix $m_{\nu_1} = 0$ to avoid numerical error, and take steps of the form $m_{\nu_2}, m_{\nu_3} \rightarrow m_{\nu_2} + \delta M_\nu/2, m_{\nu_3} + \delta M_\nu/2$ to preserve the atmospheric mass splitting.

3.5 Alcock-Paczynski effect

Power spectra, being dimensionfull quantities expressed in terms of physical distances, are usually computed from observations by assuming a “fiducial model” to convert redshifts and angular positions to 3D comoving coordinates. This fixed fiducial cosmology must be accounted for when comparing the observations to a theoretical model that predicts a different conversion to comoving coordinates. The distortion this introduces is known as the Alcock-Paczynski (A-P) effect [99]. The wavevector measured using the fiducial cosmology \mathbf{k}_{fid} is related to the true wavevector \mathbf{k} by:

$$\mathbf{k}_{\text{fid}} = \mathbf{k}_{\parallel} \alpha_{\parallel}(z) + \mathbf{k}_{\perp} \alpha_{\perp}(z). \quad (3.17)$$

where $\alpha_{\parallel}(z) = H(z)_{\text{fid}}/H(z)_{\text{true}}$ and $\alpha_{\perp}(z) = D_A(z)_{\text{true}}/D_A(z)_{\text{fid}}$. The correlation function, being a dimensionless quantity, is left unchanged by this rescaling. Thus $P_{gg}^{\text{obs}}(\mathbf{k}_{\text{fid}}, z)d^3\mathbf{k}_{\text{fid}} = P_{gg}^{\text{true}}(\mathbf{k}, z)d^3\mathbf{k}$, which gives [100]:

$$P_{gg}^{\text{obs}}(\mathbf{k}_{\text{fid}}, z) = \alpha_{\parallel}^{-1}(z)\alpha_{\perp}^{-2}(z)P_{gg}^{\text{true}}(\mathbf{k}, z). \quad (3.18)$$

This effect is included in our derivatives via the chain rule:

$$\begin{aligned} \frac{\partial P}{\partial \theta} \rightarrow \frac{\partial P}{\partial \theta} - \underbrace{\left[\frac{\partial \alpha_{\parallel}}{\partial \theta} + 2 \frac{\partial \alpha_{\perp}}{\partial \theta} \right]}_{\text{volume rescaling}} P \\ - \underbrace{\mu(1 - \mu^2) \left[\frac{\partial \alpha_{\parallel}}{\partial \theta} - \frac{\partial \alpha_{\perp}}{\partial \theta} \right]}_{-\partial_{\theta}\mu} \frac{dP}{d\mu} - k \underbrace{\left[\frac{\partial \alpha_{\parallel}}{\partial \theta} \mu^2 + \frac{\partial \alpha_{\perp}}{\partial \theta} (1 - \mu^2) \right]}_{-\partial_{\theta}k} \frac{dP}{dk}. \end{aligned} \quad (3.19)$$

3.6 BAO modeling

One of the major scientific goals of future redshift surveys is the mapping of the expansion history of the Universe using the baryon acoustic oscillation (BAO) technique [101]. The BAO features in the power spectrum form a particularly robust standard ruler, and can benefit from a technique known as “reconstruction” [102, 103], that improves the distance fidelity. Since BAO measurements are analyzed differently than many of the other parameters we’ve discussed, we forecast constraints on the distance scale differently, as we now discuss.

We hold the shape of the linear theory power spectrum fixed at a fiducial value and compute the post-reconstruction power spectrum (P_{recon}) with the “Rec-Sym” convention within the Zeldovich approximation [104]. We use the same quadratic bias model of §3.1, with fiducial bias values listed in Table 2, and approximate the reconstructed power spectrum by its first three non-zero multipoles. We include the A-P parameters, α_{\parallel} and α_{\perp} , that are defined to include a factor of the sound horizon at the drag epoch (r_d), and marginalize over a linear bias and 15 “broad band” polynomials as is usually done experimentally (e.g. refs. [105, 106]):

$$P_{\text{obs}}(k_{\perp}, k_{\parallel}; \mathbf{b}) = \alpha_{\parallel}^{-1} \alpha_{\perp}^{-2} P_{\text{recon}}(k_{\perp}/\alpha_{\perp}, k_{\parallel}/\alpha_{\parallel}, \mathbf{b}) + \sum_{n=0}^4 \sum_{m=0}^2 c_{nm} k^n \mu^{2m} \quad (3.20)$$

with $\mu = k_{\parallel}/k$ and $\mathbf{b} = (b, b_2, b_s)$. We set the fiducial value of c_{00} coefficient to the Poisson shot noise $1/\bar{n}_{\text{gal}}$, and the fiducial values of all higher order coefficients to zero.

The Fisher matrix for the reconstructed power spectrum takes the same functional form as Eq. (3.6). In our forecasts (§4.1) we interpret the marginalized errors on the tangential and radial distance measures as $\sigma_{\alpha_{\perp}} D_A/r_d$ and $\sigma_{\alpha_{\parallel}} H r_d$ respectively. In Appendix C we compare this approach to the traditional method of ref. [107].

3.7 Redshift uncertainties

An error $\sigma_z/(1+z)$ in the measurement of the redshift of an object corresponds to an error in the radial distance to said object $\sigma_{\chi} = \sigma_z d\chi/dz = [c(1+z)/H(z)]\sigma_z/(1+z)$. As a result of this error, the observed density field at Eulerian position \mathbf{x} is related to the true density field by a simple translation $\delta^{\text{obs}}(\mathbf{x}) = \delta(\mathbf{x} - \epsilon(\mathbf{x})\hat{\mathbf{n}})$, where $\hat{\mathbf{n}}$ is the direction along the line of sight, and $\epsilon(\mathbf{x})$ is a stochastic variable which we assume is drawn from a Gaussian distribution

with mean zero and standard deviation σ_χ . Thus on average, the observed density field is damped in the $\hat{\mathbf{n}}$ direction:

$$\begin{aligned} \langle \delta^{\text{obs}}(\mathbf{k}) \rangle_\epsilon &= \int d^3\mathbf{y} e^{-i\mathbf{k}\cdot\mathbf{y}} \delta(\mathbf{y}) \frac{1}{\sqrt{2\pi}\sigma_\chi} \int d[\epsilon(\mathbf{x})] e^{-\epsilon(\mathbf{x})^2/2\sigma_\chi^2} e^{-ik_\parallel\epsilon(\mathbf{x})} \\ &= e^{-k_\parallel^2\sigma_\chi^2/2} \delta(\mathbf{k}), \end{aligned} \quad (3.21)$$

which in turn damps the observed power spectrum by a factor of $\exp[-k^2\mu^2\sigma_\chi^2]$. In our forecasts we assume that the redshift uncertainty can be well characterized and ignore this extra factor in the derivatives. Therefore, at the Fisher matrix level (Eqn. (3.6)), we can in practice absorb the effect of redshift uncertainties in the noise, so that $N_0 \rightarrow N_0 \exp[k^2\mu^2\sigma_\chi^2]$.




Figure 8. $\Sigma(z)H(z)/[c(1+z)]$ as a function of redshift. If the redshift uncertainties $\sigma_z/(1+z)$ are near or above this curve, then they will suppress modes that could carry valuable cosmological information. The left y -axis shows the requirement on $\sigma_z/(1+z)$ in percent, while the right y -axis converts that to a velocity scale, in 100 km s^{-1} .

These uncertainties become increasingly important as an experiment pushes to higher redshifts. If we want to accurately reconstruct all modes with $k < k_{\text{nl}}$, we require that the uncertainty in radial length measurements is smaller than the corresponding wavelength at the non-linear scale: $\sigma_\chi < k_{\text{nl}}^{-1} = \Sigma(z)$. This in turn implies $\sigma_z/(1+z) < \Sigma(z)H(z)/[c(1+z)]$. The RHS of this inequality is plotted in Fig. 8. Obtaining small redshift uncertainties at high redshift is challenging, because the tracers are intrinsically fainter, but also because some of the broader and easier lines to measure provide intrinsically less accurate redshifts (due to asymmetric profiles, or because they are affected by winds and outflows for example) [108]. While it may be possible to use the Ly α line profile to recover the systemic velocity to $\lesssim 100 \text{ km/s}$ [109] — which would be more than sufficient for our science cases — how well this works on large samples with achievable integration times remains an open question. We delay discussing the impact of redshift uncertainties on our forecasts to §5.

4 Results

In sections 2 and 3 we outlined the surveys, samples and methods used to create our forecasts. In this section we summarize our results, estimating the constraints on base Λ CDM, curva-

ture, neutrino mass, relativistic species, primordial features, primordial non-Gaussianity, dynamical dark energy, and gravitational slip.




Figure 9. For each parameter X on the perimeter, the corresponding radius is the ratio of the uncertainty σ_X from Planck+SO+DESI (PSD) alone to the uncertainty from Planck+SO+DESI+high redshift ($z > 2$) data. The constraints on Λ CDM are calculated assuming our base model (see §3) while each extension to our base model (M_ν , N_{eff} , Ω_k) is treated as a single parameter extension. These constraints include both full shape and lensing data (see §3.3 for our conventions when combining surveys). For PSD+MegaMapper we include MegaMapper \times SO lensing, while for all other configurations DESI \times SO lensing is included. In green (orange) we show constraints for optimistic (pessimistic) foreground assumptions (§3.1).

Because we are interested in the performances of proposed future experiments in the context of current state of the art surveys, we compute the Fisher information for DESI and Euclid as well. This way all current and future surveys are studied using the same perturbation theory scheme and share the same model assumptions, allowing for a fairer comparison. In light of the existing mild tension between LSS datasets and Planck, we also present forecasted uncertainties on relevant parameters from individual surveys without any CMB external information. In doing so we obtain an estimate for the expected level of agreement (or disagreement) between disjoint LSS and CMB analyses.

As partially summarized in Fig. 9, high redshift surveys would yield significant gains in constraining power over current LSS and CMB surveys. For example, the addition of PUMA to the combination of Planck, SO and DESI tightens the constraints on ω_c , h and Ω_k by a factor of ~ 2 , while the constraints on M_ν improve by more than 60%. The addition of MegaMapper also yields to a large reduction of the uncertainties. Other extensions of the base model, such as primordial features (§4.7) and EDE (§4.8), show similar improvements, which we explore in detail in the following subsections.

4.1 Distance measures

The first results we show concern the uncertainty on the distance measures obtained from measurements of the BAO feature (§3.6). Fig. 10 shows the relative error on α_\perp and α_\parallel , which can be interpreted as relative errors on D_A/r_d and $r_d H$, for different spectroscopic and 21-cm surveys. Near-term surveys (DESI, Euclid and HIRAX) could measure the angular diameter distance to sub-percent precision all the way to $z \sim 2$, and the Hubble parameter to percent precision in the same redshift range. For 21-cm instruments we always show a shaded region encompassing optimistic and pessimistic foreground assumptions. We find that an MSE-like survey, shown in green, nicely extends the current spectroscopic program to higher redshifts,

achieving the same accuracy in the redshift range $2 \leq z \leq 4$ as DESI and Euclid at low redshift. An instrument like MegaMapper would further improve these constraints, yielding 0.5% measurements of both D_A/r_d and $r_d H$ at $z \lesssim 3.5$ and 1% measurements at $z \sim 5$. A Stage-II 21-cm experiment like PUMA-32K, shown in brown, could in principle outperform spectroscopic surveys, providing almost cosmic variance limited measurements of the Hubble parameter out to $z \sim 6$. For the measurement of distances perpendicular to the line of sight, 21-cm instruments are severely limited by the foreground contamination, but even in our pessimistic scenario PUMA-32K would reduce the uncertainties on D_A/r_d by roughly 30% compared to MegaMapper.




Figure 10. Error on the A-P parameters from the reconstructed power spectrum, which can be interpreted as relative errors on $D_A(z)/r_d$ and $r_d H(z)$ respectively. The boundaries of the shaded regions denote optimistic/pessimistic foreground assumptions for the 21-cm surveys. In the top panels we show the errorbars for the optimistic case.

We would like to stress that for 21-cm instruments it is not yet clear how to perform BAO reconstruction and therefore enhance the constraints on the distance parameters. In auto-correlation experiments new approaches have been considered [110], but no practical implementation has been proposed for interferometers. On the other hand we notice that the main purpose of BAO reconstruction is to undo large scale flows and sharpen the position of the BAO feature in configuration space. The effect of such large scale motions is smaller at high redshifts, and becomes irrelevant at $z \gtrsim 3$, after which we expect BAO reconstruction to yield diminishing returns.

Finally, it is possible to combine BAO information with the full shape analysis presented in the next few sections [111]. This would however require a model for the covariance between the two measurements, which is currently best obtained using mock catalogs. For this reason we do not include any prior on the distance measures in the forecasts discussed in the next sections.

4.2 Base Λ CDM

Historically, the standard (Λ CDM) model has been extremely successful in modeling cosmological observables, ranging from the Universe’s expansion history to the statistics of CMB anisotropies and the evolution of galaxy clustering (e.g. Fig. 3). Future LSS and CMB sur-

Configuration	$10^4 h$	$10^4 \omega_c$	$10^3 \tau$	$10^3 \Omega_k$	$10^3 N_{\text{eff}}$	M_ν [meV]
Planck	62.14	13.66	7.87	17.19	187.09	83.39
DESI	63.83	42.29	$> 10^5$	20.38	599.83	311.98
Euclid	58.02	35.55	...	17.0	476.77	245.8
MSE	49.84	21.36	...	10.79	214.09	129.18
MegaMapper	45.75	19.42	...	8.01	173.57	104.58
HIRAX	48.9/98.2	23.2/42.5	...	10.7/29.4	267.5/707.0	157.0/516.9
PUMA-5K	39.4/64.9	14.6/16.8	...	5.5/12.1	101.9/204.1	67.4/154.2
PUMA-32K	32.8/49.7	12.2/14.3	...	3.8/8.8	81.4/150.7	52.4/116.1
Planck + DESI	31.83	7.28	5.86	1.24	112.8	41.49
+ lensing	28.72	6.58	5.17	1.19	107.32	38.01
Planck + SO + DESI (PSD)	23.42	6.13	5.83	0.68	43.93	39.76
+ lensing	18.68	4.9	4.2	0.68	42.84	35.39
Planck + SO + Euclid	23.04	6.05	5.64	0.66	43.64	37.17
+ lensing	17.23	4.54	3.97	0.65	42.62	33.67
Planck + SO + MSE	17.63	4.76	5.5	0.58	41.51	31.57
+ lensing	12.47	3.41	3.46	0.57	39.27	29.2
PSD + MegaMapper	13.77	3.77	4.67	0.51	38.1	26.33
+ lensing	10.65	2.96	3.23	0.5	35.91	25.29
PSD + HIRAX	16.9/22.4	4.4/5.9	5.8/5.8	0.6/0.6	41.1/43.6	34.1/38.9
+ lensing	14.9/18.2	3.9/4.8	3.9/4.2	0.6/0.6	40.3/42.6	30.0/34.5
PSD + PUMA-5K	12.9/19.7	3.5/5.2	5.8/5.8	0.4/0.5	34.2/41.5	27.1/35.3
+ lensing	11.9/16.7	3.2/4.4	3.7/4.0	0.4/0.5	33.8/40.6	25.2/31.6
PSD + PUMA-32K	10.0/16.6	2.8/4.4	5.7/5.8	0.3/0.5	29.6/39.0	22.4/31.6
+ lensing	9.6/14.7	2.6/3.9	3.6/3.9	0.3/0.5	29.4/38.3	21.2/28.5
LiteBIRD + S4 + DESI (LSD)	21.39	5.74	2.02	0.34	24.57	34.78
+ lensing	12.86	3.41	1.93	0.34	24.02	24.38
LSD + MegaMapper	12.76	3.52	1.94	0.31	22.63	20.02
+ lensing	8.86	2.43	1.74	0.31	21.62	14.92
LSD + HIRAX	15.9/20.6	4.3/5.5	2.0/2.0	0.3/0.3	23.7/24.5	30.2/34.1
+ lensing	11.3/12.7	3.0/3.4	1.9/1.9	0.3/0.3	23.3/24.0	19.5/23.6
LSD + PUMA-5K	12.4/18.4	3.4/4.9	2.0/2.0	0.3/0.3	21.3/24.1	24.7/31.1
+ lensing	9.7/12.1	2.6/3.2	1.8/1.9	0.3/0.3	21.0/23.6	17.4/21.4
LSD + PUMA-32K	9.7/15.8	2.7/4.2	2.0/2.0	0.3/0.3	19.4/23.4	20.3/28.0
+ lensing	8.3/11.3	2.2/3.0	1.8/1.9	0.3/0.3	19.3/23.0	14.5/18.9
Everything bagel	6.95	1.94	1.69	0.24	18.36	12.24

Table 3. 1σ uncertainties on selected Λ CDM parameters and extensions from a combination of primary CMB, large-scale clustering ($P_{gg}(\mathbf{k})$), and CMB \times LSS lensing data. For all configurations that include 21-cm data, we give constraints for optimistic/pessimistic foreground assumptions. The Planck constraints in the first row are the diagonal components of the measured TT,TE,EE+lowE covariance [55]. These would be the true constraints if the Planck likelihood were a perfect Gaussian, and agree with the constraints quoted in ref. [55] to within 5% (with the exception of M_ν , where non-Gaussianities the the likelihood increase the true σ_{M_ν} to 130 meV). Everywhere else in this paper, Planck refers to the CMB prior described in §2.4. In rows 2 – 8 (LSS alone) we include a conservative $\sigma(\omega_b) = 0.0005$ prior from BBN [55]. PSD is short for Planck + SO + DESI, while LSD is short for LiteBIRD + S4 + DESI. “Everything bagel” is short for LSD + PUMA-32K (with optimistic foreground assumptions) + MegaMapper + MegaMapper \times S4 lensing, and neglects any correlations between 21-cm and LBG clustering.

veys will significantly improve over current constraints, providing the most stringent test of Λ CDM to date.

Within the next ~ 5 years, our forecasts indicate that the combination of DESI and Planck will improve upon current CMB measurements of the Hubble constant and ω_c by more than a factor of 2. By themselves, DESI and Euclid will provide similar constraints to Planck on several Λ CDM parameters, but LSS and CMB data combined are larger than the sum of their parts. In Table 3 we also show how the inclusion of CMB lensing (and its cross-correlation with galaxies) improves the constraints. For Planck lensing and current or upcoming surveys like DESI and Euclid, the addition of lensing does not yield significant improvements. Given the time frame of future LSS surveys, we expect to (in principle) be able to combine them with Planck+SO+DESI (PSD) data. The addition of MegaMapper to PSD improves the constraint on the Hubble constant by almost a factor of 2, and by $\sim 40\%$ on ω_c . We find that cross-correlations with CMB lensing maps return larger gains for a MegaMapper-like survey, presumably due to the increased overlap with the CMB lensing kernel. For 21-cm surveys, foregrounds will likely make the cross correlation with CMB lensing maps not possible. Thus, in Table 3, the “+lensing” rows below any PSD+21 cm surveys only include the cross-correlation of DESI with SO (or S4) lensing maps. Nonetheless PSD+PUMA-5K or -32K could yield comparable constraints to PSD+MegaMapper on most Λ CDM parameters. In the future, density reconstruction methods could help 21-cm surveys to recover the modes lost in the foreground wedge, allowing for a non-zero correlation with CMB lensing maps [112, 113].

On a similar timescale to future LSS surveys, CMB experiments like LiteBIRD and CMB-S4 will also be taking data. It is therefore interesting to consider their combination with DESI as the baseline, which we dub LSD. As shown in Table 3, clustering data at high redshifts enable significant improvements in Λ CDM constraints over a nearly cosmic variance limited CMB survey. The uncertainties on h and ω_c from LSD only are $> 50\%$ larger than those from LSD+MegaMapper. Improvements from cross-correlations with CMB lensing are also of the same order. We note that constraints from LSD+PUMA-32K are comparable to those from LSD+MegaMapper, despite 21-cm surveys having a lower effective shot noise. This is ultimately due to degeneracies between the brightness temperature T_b , A_s and the linear bias b , which in turn worsen the constraints on any cosmological parameters that the latter are degenerate with.

4.3 Structure growth

Future high-redshift surveys will enable 1% measurements of the amplitude of structure growth $\sigma_8(z)$ out to $z \sim 5$. We consider two different methodologies for structure growth measurements: one in which σ_8 is treated as a derived parameter from Λ CDM, and another, more cosmology-agnostic approach in which the shape of the power spectrum is held fixed.

In the derived-from- Λ CDM procedure we calculate the Fisher matrix from full shape clustering data with basis $\{\theta_i\} = \{\Lambda\text{CDM}, b, b_2, b_s, N_0, N_2, N_4, \alpha_0, \alpha_2, \alpha_4\}$ in each redshift bin. Our choice of Λ CDM parameters is listed in Table 2. We then transform to a basis that includes $\sigma_8(z)$, where z is the central redshift of some redshift bin. The Fisher matrix transforms as a two-form under a change of basis: $\mathbf{F} \rightarrow \mathbf{R}^T \mathbf{F} \mathbf{R}$, where

$$\mathbf{R}^{-1} = \begin{pmatrix} \frac{\partial \sigma_8(z)}{\partial \theta_1} & \frac{\partial \sigma_8(z)}{\partial \theta_2} & \dots & \frac{\partial \sigma_8(z)}{\partial \theta_n} \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{pmatrix}. \quad (4.1)$$

We then add a conservative $\sigma(\omega_b) = 0.0005$ [55] prior from BBN, and invert this matrix to yield the marginalized constraint on $\sigma_8(z)$. Thus the measurements of $\sigma_8(z)$ are local, derived only from data within the redshift bin of interest, apart from the shared ω_b prior. As shown in the left panel of Fig. 11, DESI and Euclid will measure $\sigma_8(z)$ to 5 – 10% across the range $0.6 < z < 2$, while MSE and MegaMapper will extend these measurements out to $z = 5$ with less than 5% error. Due to a degeneracy between σ_8 and T_b , the constraints from HIRAX are more modest than those from DESI or Euclid. Since σ_8 and T_b are perfectly degenerate in linear theory, 21-cm measurements of σ_8 heavily depend on how accurately one can model and measure non-linear effects [20]. While PUMA also suffers from the same $\sigma_8 - T_b$ degeneracy, it has a much lower shot noise than HIRAX, and is thus more sensitive to non-linear scales that differentiate the effects of the two parameters, enabling 5 – 10% measurements across its redshift range. We find that PUMA-32K (with a BBN prior on ω_b) can measure $T_b(z)$ to the same relative accuracy as $\sigma_8(z)$, or 5 – 10% depending on the redshift bin. Thus a $\approx 1\%$ prior on $T_b(z)$ would be required for 21-cm measurements of $\sigma_8(z)$ to be competitive with MegaMapper, which is likely beyond the reach of astrophysical determinations from high-column-density systems.




Figure 11. *Left:* Errors on $\sigma_8(z)$ when derived from Λ CDM, assuming a $\sigma(\omega_b) = 0.0005$ prior from BBN. *Right:* Errors on $\sigma_8(z)$ from full shape data when calculated using the fixed shape procedure, without any external priors. The errorbars with smaller caps include cross-correlations with SO lensing, but do not include any information from the convergence auto-spectrum. All 21-cm measurements assume optimistic foregrounds. The gray bands extend from 1 ± 0.01 .

In Fig. 11 we also show the improvement from adding $\text{SO} \times \text{LSS}$ lensing to full shape data. When adding lensing to this figure we only include the cross-correlation, and not the convergence power spectrum. This ensures that the constraints on $\sigma_8(z)$ remain local apart from the ω_b prior. Thus these constraints make no assumptions about the expansion history outside of the redshift bin of interest. As was true in the previous section, lensing cross-correlations offer small improvements for DESI and Euclid, while for MegaMapper, the improvement can be more than a factor of 2. With lensing cross-correlations, MegaMapper can measure $\sigma_8(z)$ with a 1% error across its entire redshift range.

While the derived-from- Λ CDM constraints are only valid within the Λ CDM framework, fixed shape constraints assume very little about the underlying cosmology. For any cosmology whose free parameters include an amplitude of the matter power spectrum, one can obtain a fixed shape measurement by fitting to the amplitude with all other parameters fixed at some fiducial value. Thus fixed-shape forecasts give a sense of a “best-case” scenario, since the error on the amplitude from the fixed shape procedure is always smaller than the error derived from

the full cosmological model. In our case we fix $\{h, \omega_c, \omega_b, n_s, \tau\}$ at their fiducial Λ CDM values, and treat the amplitude of the matter power spectrum as an independent variable in each redshift bin. In each bin we marginalize over the nuisance terms listed in Table 2. With this procedure, we find that DESI and Euclid yield 2% measurements across $0.6 < z < 2$, while MSE and MegaMapper extend these measurements out to $z = 5$ with comparable uncertainty. When cross-correlations with SO lensing are included, these errors decrease to sub-percent values across the entire redshift range covered by spectroscopic surveys ($0.6 < z < 5$). Again, we see that 21-cm measurements suffer from a $\sigma_8 - T_b$ degeneracy, resulting in much more modest $\mathcal{O}(10\%)$ constraints.

We find that fixed shape measurements from spectroscopic surveys have 2 to 5 times smaller error than when derived from Λ CDM (without lensing cross-correlations). For 21-cm surveys, however, the two methods yield similar constraints, presumably due to the $\sigma_8 - T_b$ degeneracy dominating over any other degeneracy between σ_8 and the remaining Λ CDM parameters, such as τ .

4.4 Massive neutrinos

One of the best-motivated extensions beyond the Λ CDM model described above is allowing the neutrino masses to be non-minimal [114]. Measurements of solar and atmospheric neutrino oscillations probe $m_{\nu_2}^2 - m_{\nu_1}^2$ and $m_{\nu_2}^2 - m_{\nu_3}^2$ respectively, which together imply a minimum value for the sum of the neutrino masses $M_\nu = \sum_i m_{\nu_i} > 60$ meV. Oscillation measurements have yet to determine if neutrinos follow a normal ($m_{\nu_2} < m_{\nu_3}$) or inverted ($m_{\nu_2} > m_{\nu_3}$) mass hierarchy, which is one of the major goals of upcoming oscillation measurements and neutrino-less double beta decay experiments. Cosmological constraints of M_ν are complementary to these experiments. CMB and BAO data currently provide our tightest upper limits on M_ν [115], and future surveys should tighten these constraints significantly.

Neutrinos are relativistic at early times, washing out their primordial fluctuations at length scales smaller than the horizon. However, once the neutrino fluid cools and becomes non-relativistic (occurring at a scale factor a_{nr}), so that its energy density is dominated by the neutrino mass M_ν , neutrinos cluster in the same way as the CDM+baryon fluid. The comoving length scale associated with this transition is $k_{\text{nr}} = a_{\text{nr}}H(a_{\text{nr}}) \sim 10^{-3}h\text{Mpc}^{-1}$ [14]. Thus, another way of saying the above is that neutrinos only cluster at large scales ($k \ll k_{\text{nr}}$). At small scales, the neutrino energy density is nearly spatially homogeneous, contributing to the Hubble drag associated with the evolution of the CDM+baryon fluid, but not to the gravitational potential, which drives the clustering. Thus a nonzero neutrino mass slows the growth of structure, damping the matter power spectrum at scales smaller than the horizon at a_{nr} .

Since this effect is largest at small scales and at late times, high redshift LSS surveys – which will measure high k modes in the late-time Universe – are ideally situated for M_ν measurements. As shown in Table 3, we see that $\sigma(M_\nu)$ drops dramatically with the addition of high redshift data: the constraint from PSD is 35 (40) meV with(out) lensing, which drops to 25 (26) meV for PSD+MegaMapper, or 21 (22) meV for PSD+PUMA-32K. As shown in Fig. 12, these results depend significantly on the k_{max} used in the forecast. The constraints change by nearly a factor of two for a LSD+MegaMapper-like survey between $0.2 < k_{\text{max}}/k_{\text{nl}} < 1.0$. Likewise, the 21-cm foreground assumptions have an equally significant effect on the forecasted neutrino constraints, impacting the results by up to 50% for the highest values of k_{max} for PSD+PUMA-32K⁷. However, even with a more conservative

⁷For parameters with a physical priors, *e.g.* $M_\nu > 0$, the Fisher matrix approach does not necessarily

cutoff of $k_{\max}/k_{\text{nl}} = 0.5$, and with mildly conservative foreground assumptions, a next generation CMB survey coupled to either MegaMapper or PUMA-32K could comfortably achieve $\sigma(M_\nu) \lesssim 20$ meV.




Figure 12. Constraints on M_ν as a function of k_{\max}/k_{nl} from the primary CMB and full shape data. The curves that are labeled “(with lensing)” include information from DESI×S4 lensing ($\kappa\kappa$, κg , gg). In both plots we set $k_{\max} = k_{\text{nl}}$ for DESI, and only vary k_{\max}/k_{nl} for the high redshift survey of interest. *Left:* Constraints for a MegaMapper-like survey that observes a redshift-independent fraction \bar{n}/n_{LGB} of LBGs from $2 < z < 5$ with $f_{\text{sky}} = 0.34$. *Right:* Constraints for PUMA. The edges of the shaded regions represent optimistic and pessimistic foreground assumptions.

Currently in the literature there is a significant bifurcation in M_ν forecasts for the upcoming Euclid satellite. Ref. [64] found 34 meV for Planck+SO+Euclid+Euclid×SO lensing using a PT-based method in which nine nuisance terms were marginalized over, including four bias parameters (b , b_2 , b_s , b_3), three counterterms (α_0 , α_2 , α_4) and two stochastic contributions (N_0 , N_2). When we fix N_4 and match their $k_{\max} = 0.2 h \text{ Mpc}^{-1}$ cut we find 40 meV for Planck+SO+Euclid+Euclid×SO lensing. A $\sim 15\%$ discrepancy is entirely reasonable given that their Euclid sample spans a broader redshift range ($0.6 < z < 2.1$) than ours ($0.9 < z < 1.8$), uses slightly different number densities and biases, and also marginalizes over a b_3 term. By contrast, ref. [63] finds 17 meV for Planck+Euclid+lensing (PS only) using a PT-based approach in which seven nuisance parameters were marginalized over, including three bias terms, three counterterms, and the shot noise. When we fix N_2 and N_4 and set $k_{\max} = 0.5 h \text{ Mpc}^{-1}$, so that our set of nuisance terms and k_{\max} match ref. [63], we find 30 meV for Planck+Euclid+lensing+BAO. If we fix τ our constraint drops to 20 meV, in good agreement with ref. [63]. The residual difference could be due to the fact that the authors in [63] employ a full MCMC approach as opposed to the Fisher matrix presented in this work.

Similar results for the cross-correlation of galaxies with CMB lensing were also recently obtained by [65], who consider a non-linear bias expansion similar to ours. Our results are largely consistent with recent forecasts for upcoming high redshift surveys. Ref. [9] found 26

return trustworthy error bars if the fiducial value chosen is close to the physical boundary. For this reason we caution to reader about the interpretation of our results when $\sigma(M_\nu)$ is large.

meV for PSD+MegaMapper⁸, while we find 25 meV. Recent PUMA forecasts [39] show M_ν improving from 38 meV with Planck+LSST+DESI to 20 meV with CMB-S4+PUMA-32K, while we get 40 meV for Planck+SO+DESI and 20 meV for LSD+PUMA-32K.

We note that the inclusion of a future cosmic-variance limited measurement of τ from large scale CMB polarization such as from the proposed LiteBIRD satellite provides significant improvements on M_ν by removing the $\tau - A_s$ degeneracy present in the primary CMB. The very large-scale polarization signal is challenging to measure from the ground, potentially requiring a future satellite mission⁹. However, small scale information in both the CMB temperature (through fluctuations in the kinematic Sunyaev-Zel’dovich effect [117, 118]), or the 21cm signal from reionization [119, 120] can potentially constrain τ to a similar level as LiteBIRD ($\sigma(\tau) \approx 0.002$), if astrophysical systematics can be kept under control. Including such a prior on τ reduces the PSD+MegaMapper neutrino mass constraint from 25 meV to 16 meV (including lensing) while for PSD+PUMA-32K it drops from 21 meV to 16 meV. As shown in Table 3 and Fig. 20, these improvements are similar to those obtained by upgrading from PSD to LSD, suggesting that the improvement from these future CMB data is almost entirely dominated by the better τ measurement.

Finally we note that with a LiteBIRD-like prior on τ and $k_{\text{max}}/k_{\text{nl}} = 1$, MegaMapper (which corresponds to $\bar{n}/n_{\text{LBG}} \approx 0.5$) is primarily limited by its shot noise, and that small improvements are possible with a higher LBG target efficiency. Increasing the number density to $\bar{n}/n_{\text{LBG}} = 1$ decreases the LSD + MegaMapper (without lensing) constraint from 20 to 18 meV.

4.5 Light relics

Measurements of the total amount of radiation in the Universe are a powerful way to search for new physics. Any light particle that was once in thermal equilibrium with the Standard Model, and hence had a cosmological abundance, leads to well known modifications in the power spectrum of the CMB and LSS. These effects are commonly parameterized through N_{eff} , defined such that

$$\rho_r = \frac{\pi^2}{15} \left[1 + \frac{7}{8} \left(\frac{4}{11} \right)^{4/3} N_{\text{eff}} \right] T_\gamma^4 \quad (4.2)$$

during the radiation dominated era, with $N_{\text{eff}} = 3.046$ [121] in the Standard Model (or $N_{\text{eff}} = 3$ in the limit of instantaneous neutrino decoupling). The presence of new light relics would shift this value to $N_{\text{eff}} = 3.046 + \Delta N_{\text{eff}}$, where ΔN_{eff} depends on the number of independent spin states and decoupling temperature(s). By considering decoupling temperatures that are significantly larger than the top quark mass, one obtains lower limits to ΔN_{eff} for broad classes of particles: $\Delta N_{\text{eff}} > 0.027$ for scalars, 0.047 for Weyl fermions, and 0.054 for vector bosons [14]. In many examples of physics beyond the Standard Model, these new states are too weakly coupled to be probed by facilities on Earth, therefore highlighting the synergy between cosmological probes and traditional particle physics experiments.

In our forecast we keep all other parameters fixed while varying N_{eff} , therefore the presence of extra radiation produces a change in the position of the BAO peaks and a

⁸This includes BAO only and no lensing, uses a different k_{max} , and does not marginalize over non-linear terms, so the comparison is not entirely fair.

⁹The CLASS experiment [116] aims to measure these very large-scale polarization fluctuations from the ground, and depending on performance, it may lead to competitive constraints on τ compared to a future space mission.

damping of the matter power spectrum on small scales [122], as shown in Fig. 7. The density perturbations of the new degrees of freedom also produce a phase shift in the oscillation scale of the baryon-photon fluid before recombination [122]. In the CMB power spectrum, the small-scale suppression due to extra relativistic species is largely determined by a shift in the damping scale k_d , which is sourced by a change in the early time Hubble parameter $H(z)$ through the change in ρ_r . As such, N_{eff} is rather degenerate with the primordial helium abundance Y_p in CMB data, and marginalizing over Y_p can have a significant effect on N_{eff} constraints, weakening them by up to a factor of two [123]. Note that this degradation is present despite the N_{eff} -induced phase shift in the CMB peaks location which is independent of Y_p and therefore partly helps break this degeneracy, suggesting that the phase-shift alone is not playing a major role.

By contrast, the growth of dark matter density fluctuations on scales probed by LSS experiments is sensitive to the modified expansion history introduced by new free streaming particles, and it is not affected by Silk damping. It is therefore largely insensitive to Y_p , which is thus not included in our analysis. The BAO phase shift is also not the major source of information on N_{eff} in galaxy clustering data [124], whose constraining power is mostly driven by the change in the full shape of the power spectrum.




Figure 13. N_{eff} constraints from Planck+SO (primary CMB) +DESI+PUMA-NK (full shape) for variable dish number N and z_{max} . No lensing data is included in this figure. In the right plot we fix $k_{\text{max}}/k_{\text{nl}} = 1$, while in the left plot we fix $N = 32$. For reference, the constraint from Planck+SO (primary CMB) +DESI (full shape) is 0.044.

In our forecasts we marginalize over Λ CDM and the nuisance terms listed in Table 2. Our results are summarized in Table 3 and Fig. 13. Near term LSS surveys (DESI, Euclid, MSE) will improve over the SO goal of $\sigma(N_{\text{eff}}) = 0.05$ by only $\sim 15\%$. The addition of high redshift data reduces this uncertainty to 0.03, a 40% improvement over Stage-III CMB measurements alone. With a Stage-IV CMB survey these constraints drop to 0.019, or 0.018 when a combination of high redshift LBG and 21-cm data are included. A measurement with uncertainty $\sigma(N_{\text{eff}}) = 0.018$ could detect light scalars (vectors) at 1.5σ (3σ) significance, regardless of their freeze-out temperature, and would be capable of detecting (or ruling out) all light relics with a freeze-out temperature smaller than the QCD phase transition (~ 0.2 GeV, or $\Delta N_{\text{eff}} \sim 0.1$) at $> 5\sigma$ significance [14].

As shown in Fig. 13 the constraints from a PUMA-NK survey plateau with detector number N at around $N = 32$, suggesting that PUMA-32K constraints are limited by the foreground removal of modes inside the wedge. From Table 3, we see that both MegaMapper

and PUMA-32K achieve similar constraints (~ 0.02 with CMB-S4). Thus MegaMapper is also sufficient for achieving a nearly cosmic-variance limited N_{eff} measurement from LSS. As one would expect, from Fig. 13 we see that these constraints significantly depend on the k_{max} assumed in the forecast, changing at the $\sim 25\%$ level between $0.1 < k_{\text{max}}/k_{\text{nl}} < 1$. From Fig. 13 we also see that decreasing PUMA’s redshift range to $2 < z < 4$, which halves the survey volume, only increases the N_{eff} constraint by $\sim 10\%$, which implies that most of the information is coming from the breaking of the degeneracies between cosmological parameters obtained by combining CMB with LSS data.

Our results are in good agreement with ref. [39], which show N_{eff} constraints dropping from 0.026 with CMB-S4 to 0.013 with CMB-S4 + PUMA-32K. Without lensing, our constraints are 0.019 for LSD + PUMA-32K. This apparent discrepancy can be entirely explained by differences in nuisance parameters. When non-linear nuisance parameters ($b_2, b_s, N_2, N_4, \alpha_2, \alpha_4$) are held fixed, we find $\sigma(N_{\text{eff}}) \sim 0.012$, which is in excellent agreement with ref. [39]. This suggests that improved models of galaxy formation, which could yield tight priors on non-linear nuisance parameters, would significantly improve a LSS measurement of N_{eff} , with the “extra information” primarily coming from small scales.

As previously mentioned, marginalizing over Y_p can significantly degrade the constraints on N_{eff} from CMB surveys. For example, the constraints from Planck increase from 0.19 to 0.30 with Y_p marginalized over. However, when clustering data from high redshift surveys are included, we find that the resulting constraint on N_{eff} is largely immune to marginalization over Y_p . For example, the constraint from Planck + DESI + PUMA-32K only increases from 0.046 to 0.052. This is comparable to the SO goal of 0.05, suggesting that future high-redshift LSS surveys will have a similar constraining power to a Stage-III CMB survey, but with the added benefit of being immune to Y_p .

4.6 Primordial non-Gaussianity

In the simplest inflationary models one expands the action of a single scalar field minimally coupled to gravity about a spatially homogeneous solution. The leading order corrections to the action are quadratic in the small fluctuations about this solution. Thus these fluctuations are drawn from a Gaussian distribution to lowest order. Higher order terms in the action, arising from interactions of the inflaton with itself or other fields, result in trace amounts of Primordial non-Gaussianity (PNG). Of particular interest are local PNG, which vanish in the squeezed limit for slow-roll single field models of inflation [125, 126].

The non-Gaussian contributions to the primordial potential Φ are commonly modeled as a perturbative expansion in a Gaussian random field ϕ [127]:

$$\Phi(\mathbf{x}) = \phi(\mathbf{x}) + f_{\text{NL}} \int d^3y d^3z W(\mathbf{y}, \mathbf{z}) \phi(\mathbf{x} + \mathbf{y}) \phi(\mathbf{x} + \mathbf{z}) + \mathcal{O}(\phi^3), \quad (4.3)$$

where $W(\mathbf{y}, \mathbf{z})$ is some kernel and f_{NL} controls the amplitude of the non-Gaussianity. In this paper we specify to the local type of PNG in which $W(\mathbf{y}, \mathbf{z}) = \delta^D(\mathbf{y})\delta^D(\mathbf{z})$. Currently the tightest constraints on the amplitude of local PNG $f_{\text{NL}}^{\text{Loc}} = -0.9 \pm 5.1$ come from the Planck satellite [115], which are within a factor of 3 of a cosmic-variance limited CMB survey [26]. LSS surveys are not yet competitive with the CMB, with current constraints around $\sigma(f_{\text{NL}}^{\text{Loc}}) \sim 25$ [128]. Pushing beyond $\sigma(f_{\text{NL}}^{\text{Loc}}) \sim 1$ requires the addition of future 3D LSS data, which has access to significantly more modes than the 2D CMB on large scales.

In addition to modifying the Universe’s initial conditions (e.g. producing a non-zero bispectrum) PNG modifies the non-linear evolution of structure growth (producing new




Figure 14. Constraints on local primordial non-Gaussianity as a function of k_{\min} and k_{\max} , using the full shape power spectrum without any CMB priors. For SPHEREx we set $\sigma_z/(1+z) = 0.05$, while for all other surveys we set $\sigma_z = 0$. The vertical lines are the lowest possible k_{\min} for each survey: $2\pi/V^{1/3}$, where V is the volume of the survey (including sky coverage).

mode-couplings), which can in principle be modeled using standard perturbative techniques. However, the development of a PT-based code which self-consistently models non-linearities (in redshift space) in the presence of PNG is beyond the scope of this paper. We instead focus on PNG’s effect on the linear bias, and restrict our forecasts to $k_{\max} < 0.06 h\text{Mpc}^{-1}$ where linear theory is a good approximation for the high redshifts we consider. As seen in Eq. 4.3, PNG couples short wavelength modes of Φ to long wavelength modes of ϕ . This coupling induces an additional scale-dependent bias ($b \rightarrow b + \Delta b$) which takes the form [129]:

$$\Delta b \equiv \frac{3\Omega_m H_0^2}{2k^2 D(z) T(k)} b_\phi f_{\text{NL}}^{\text{Loc}} = \frac{3\Omega_m \delta_c H_0^2}{k^2 D(z) T(k)} (b-1) f_{\text{NL}}^{\text{Loc}}, \quad (4.4)$$

where $D(z)$ is normalized to a in the matter dominated era and $\delta_c = 1.686$. In the above equation we made the assumption that $b_\phi = 2\delta_c(b-1)$ [130], which does not necessarily hold for the galaxies that will be observed by current and future LSS surveys [131–133], and it cautions against combining the constraints from different instruments.

Shown in Fig. 14 are our forecasted constraints on $f_{\text{NL}}^{\text{Loc}}$ from a variety of upcoming or proposed LSS surveys, using the power spectrum only. Since all information about $f_{\text{NL}}^{\text{Loc}}$ comes from low k where linear theory is an excellent approximation, in our forecasts we only marginalize over ΛCDM , b and T_b (for 21-cm surveys). For all surveys we include a prior on ΛCDM from Planck. Even with just the power spectrum, we see that near-term LSS surveys (such as Euclid) have the capability of achieving similar constraints on PNG as Planck, while a future high-redshift LBG survey has the potential to achieve $\sigma(f_{\text{NL}}^{\text{Loc}}) \sim 0.8$, more than a factor of 2 improvement over a cosmic variance limited CMB survey. Similar constraints from the power spectrum should be obtainable with the SPHEREx mission [26].

Here we have only focused on local PNG from the power spectrum, making use of the $f_{\text{NL}}^{\text{Loc}}$ induced scale-dependence of linear bias. We note that adding bispectrum or higher point function information can potentially help [133–135]. Extracting bispectrum information will be crucial for constraining shapes of PNG beyond the local type considered here (such as

equilateral or orthogonal, for example), where the effects on the power spectrum are largely degenerate with non linear evolution [136].

Finally we note that the signatures of a non-zero $f_{\text{NL}}^{\text{Loc}}$ appear on very large scales, which are often most affected by systematics (especially angular measurements). In Section 5 we explore deprojecting angular modes by introducing a minimum k_{\parallel} cut. Alternatively, one can use the fact that most galaxy imaging systematics are uncorrelated with CMB lensing maps, and measure the scale dependent bias by cross-correlating the galaxy positions with CMB lensing. Ref. [137] found that $\sigma(f_{\text{NL}}^{\text{Loc}}) \lesssim 1$ can be achieved by future surveys using this technique, thus providing an important independent check on systematics.

4.7 Primordial features

Current observations of the CMB are consistent with the primordial potential having a nearly scale-invariant power spectrum. This is a generic prediction of single-field slow-roll inflation, sourced by the (near) flatness of the inflationary potential [138]. However, attempts to connect this simple picture of inflation with more fundamental ultra-violet scenarios often result in deviation from scale invariance on some scales [1]. In this paper we focus on a subclass of these features (sharp features), which are commonly modeled as a step in the inflationary potential [139]. This step produces a scale-dependent oscillatory feature in the primordial power spectrum, which we model as a sinusoidal modulation of the linear power spectrum: $P_m(k) \rightarrow [1 + A_{\text{lin}} \sin(\omega_{\text{lin}} k + \phi_{\text{lin}})] P_m(k)$ [11, 140].




Figure 15. Constraints on A_{lin} for various fiducial values of ω_{lin} , assuming $\phi_{\text{lin}} = \pi/2$. For all surveys we include a Planck+SO prior on ΛCDM . The slight “bump” in the constraints at $\omega_{\text{lin}} \sim 100 h^{-1}\text{Mpc}$ is due to the degeneracy with the BAO frequency.

In our forecasts we fix both ω_{lin} and ϕ_{lin} , and find the corresponding uncertainty on A_{lin} after marginalizing over the parameters in Table 2. Our results are shown in Fig. 15 for $10 h^{-1}\text{Mpc} < \omega_{\text{lin}} < 300 h^{-1}\text{Mpc}$, assuming $k_{\text{max}} = k_{\text{nl}}^{\text{10}}$. We find that near-term

¹⁰Since the high frequency oscillatory feature induced by primordial features is immune to small-scale nonlinearities [140], one could in principle extend the k_{max} used in an analysis beyond the nonlinear scale. In

spectroscopic surveys will be capable of achieving $\sigma(A_{\text{lin}}) \sim 2 \times 10^{-3}$, while HIRAX has the capability to improve upon these constraints by a factor of two. Order of magnitude improvement are possible by extending out to high redshift, with an optimistic configuration of PUMA-32K achieving $\sigma(A_{\text{lin}}) \sim 2 \times 10^{-4}$.

These results are largely consistent with ref. [140], who found $\sigma(A_{\text{lin}}) \simeq 2 \times 10^{-3}, 1.5 \times 10^{-3}$, and 10^{-4} for DESI, Euclid, and a cosmic variance limited LSS survey (out to $z = 6$) respectively. Our results for DESI and Euclid are within $\sim 25\%$ of these values, which is entirely reasonable given that ref. [140] uses a different forecasting approach in which $P_{gg}^w(k)/P_{gg}^{\text{nw}}(k)$ is the observable of interest. This suggests that PUMA-32K (with optimistic foreground assumptions) is within a factor of two of a cosmic variance limited survey. We note that the slight ‘‘bump’’ in our constraints around $\omega_{\text{lin}} \sim 100 h^{-1}\text{Mpc}$ is smaller than the bump found by ref. [140], which is due to a degeneracy with the BAO frequency. This suggests that the full shape information is sufficient to break the degeneracy between primordial features and the BAO feature.

4.8 Dynamical dark energy

Models of dark energy are typically invoked to explain the observed accelerated expansion of the Universe at the present epoch. Current observations are consistent with the existence of a cosmological constant Λ . In this section we quantify the sensitivity of future surveys to departures from the Λ -model, using both a ‘‘model-independent’’ approach¹¹, and through constraints on a specific model of dynamical dark energy.




Figure 16. The absolute error on the dark energy fraction (Ω_{DE}) achievable by current and future LSS surveys.

Over the redshift range covered by our observations (typically up to $z \sim 5 - 6$), direct expansion and growth measurements can constrain any dynamical dark energy model by quantifying the fraction of dark energy in a number of redshift bins as we now explain. At redshifts $z < 6$ the Universe is dominated by nonrelativistic matter and dark energy, so that the first Friedmann equation is approximately given by $\Omega_{\text{DE}}(z) + \Omega_m(z) = 1$, where $\Omega_m(z) = \rho_m(z)/\rho_{\text{crit}}(z) = \omega_m(1+z)^3/h^2(z)$ and $h(z)$ is defined through $H(z) = 100 h(z)$ km/s/Mpc. We estimate the error on the dark energy fraction via standard error propagation, and neglect the covariance between $h(z)$ and ω_m , so that $\sigma_{\Omega_{\text{DE}}(z)}^2 = \Omega_m^2(z)(4[\sigma_{h(z)}/h(z)]^2 + [\sigma_{\omega_m}/\omega_m]^2)$.

practice, however, we find that both the MegaMapper and PUMA-32K constraints improve by at most 2% when k_{max} is increased from k_{nl} to $1.4 k_{\text{nl}}$. For MegaMapper this is mostly due to the large value of the shot noise, which limits the ability of reconstructing oscillatory features on small scales.

¹¹Ref. [141] performs a similar analysis to estimate constraints on $\Omega_{\text{DE}}(z)$, focusing on the redshift interval $1.5 < z < 2.5$.

We take the relative error on $h(z)$ from the radial BAO measurements ($\sigma_{h(z)}/h(z) = \sigma_{\alpha_{\parallel}(z)}$), which are described in sections 3.6 and 4.1. We assume that the measurement of ω_m is driven primarily from CMB measurements, and take $\sigma_{\omega_m} = 2.4 \times 10^{-4}$ for all surveys, motivated by the constraints achievable by LSD + MegaMapper (see Table 3). Our results are shown in Fig. 16. We find that MegaMapper will be able to constrain the dark energy fraction to better than 2% across its redshift range, while PUMA-32K could improve these measurements to the sub-percent level.

While the simplest models have Ω_{DE} drop rapidly at higher redshift, so that dark energy is a late-time phenomenon, there is little guidance from theory as to how ρ_{DE} should evolve, and the presence of an inflationary epoch in the very early Universe calls into question whether there should be only one late-time epoch of dark energy domination. Recently, early dark energy (EDE) has been invoked to decrease the CMB sound horizon via an increase in the expansion rate before recombination. Doing so would increase the inferred value of H_0 from CMB observations, making EDE a possible candidate for resolving the H_0 tension [142]. Current constraints from a combination of CMB, SNe and LSS data tend to disfavor this explanation [4–7], with much of the constraint coming from EDE’s effects on the growth of large-scale structure.




Figure 17. Change in the linear matter power spectrum (at $z = 2$) when including Early Dark Energy (EDE) for various values of f_{EDE} and $1.5 < \log_{10}(z_c) < 6.5$. $\theta_i = 2.83$, and all standard cosmological parameters are held at their fiducial ΛCDM values. For $z_c \lesssim 30$, the ratio $P_{\text{EDE}}^{\text{lin}}(k)/P_{\Lambda\text{CDM}}^{\text{lin}}(k)$ is roughly given by $1 - f_{\text{EDE}}$.

This effect is much more general than the above-mentioned example however, and future LSS surveys can surprisingly put very tight constraints on EDE models at redshift much higher than the range covered by the observations, thanks to the exquisite precision with which they measure the shape of the power spectrum. Since the growth of structure is damped by Hubble expansion, long periods where unclustered species dominate the expansion lead to suppression of large-scale structure that can be detected by comparing early- and late-time measures of the fluctuation amplitude. A precise measurement of the power spectrum shape can also be used to place constraints on (or detect) short-term deviations from matter or radiation domination, since such periods will change the shape of the power spectrum due

to differential growth of modes. An accurate measurement of the power spectrum constrains deviations in the expansion history over a broad range of redshifts [8].

We give an illustrative example in Fig. 17, which considers the EDE scenario of ref. [143], in which a light scalar ϕ with potential $V(\phi) \propto [1 - \cos(\phi/f)]^3$ is initially displaced from its minimum. The field is parameterized by its maximum energy contribution, $f_{\text{EDE}} \equiv \max[\rho_{\text{EDE}}(z)/\rho_{\text{tot}}(z)]$, which occurs at redshift z_c , and the initial value of the field $\theta_i \equiv \phi_i/f$. Values of $z_c \gtrsim z_*$, where $z_* \sim 1100$ is recombination, would help to relieve the Hubble tension, but we show that future large-scale structure surveys can probe even small values of f_{EDE} over a very broad range of redshifts. Using the publicly available code `CLASS_EDE` [4] to produce the initial power spectra and `velocileptors` to treat the damping of features from non-linear evolution and biasing [8], we forecast constraints on f_{EDE} for values of z_c ranging from ~ 30 to 10^6 . In all forecasts z_c and $\theta_i = 2.83$ are held fixed, and we marginalize over the parameters listed in Table 2. The results are shown in Fig. 18.




Figure 18. Constraints on the maximum amplitude of early dark energy (f_{EDE}) as a function of the time at which EDE peaks z_c , assuming $\theta_i = 2.83$. We include a Planck+SO prior on ΛCDM for all experiments. In the left panel we show constraints from full shape (FS) measurements only, while in the right panel we include a prior on ΛCDM and the nuisance parameters in Table 2 from SO lensing and cross-correlations with the respective galaxy surveys, following the convention in §3.3.

We see that near-term spectroscopic surveys are capable of constraining the amplitude of early dark energy at the few percent level out to redshift 10^4 , while PUMA would be capable of achieving subpercent level constraints out to redshift 10^5 . These are order of magnitude improvements over the $\mathcal{O}(10\%)$ constraints found by refs. [4, 144] from a combination of current BOSS and Planck data. As shown in Fig. 17, EDE that peaks prior to recombination produces a phase shift in the BAO feature. From Fig. 17 we see that this effect peaks at around $z_c \sim 10^{3.5}$, which sensibly corresponds to the minimum in Fig. 18. However, EDE that peaks after recombination induces no phase shift in the BAO signature. In this case, EDE damps the linear power spectrum in a scale-independent way (for $k > 0.01 h\text{Mpc}^{-1}$) that is nearly degenerate with A_s , b and T_b . Thus all of the information about post-recombination EDE comes from low k . This puts 21-cm observations at a disadvantage, as low k modes are excluded by the foreground wedge. We see this effect in Fig. 18. PUMA-32K offers no improvement over DESI for $z_c < z_*$, while MegaMapper improves over DESI by more than a factor of two. We note that constraints in the $z_c < z_*$ regime benefit significantly from the addition of lensing cross-correlations (MegaMapper constraints decrease by more than a factor of ~ 2), which soften the $f_{\text{EDE}} - b$ degeneracy.

4.9 Modified gravity

Modifications to GR have been proposed as an alternative explanation for the observed accelerated expansion of the Universe at low redshift [1, 17–19]. The phenomenology of such models can be very rich, but a common ingredient is that the responses of the Newtonian potential Ψ and spatial curvature Φ to matter and energy (which are equal in the low redshift Universe in GR) can differ.

To estimate the sensitivity of future surveys to deviations from GR we examine the gravitational slip parameter $\gamma = \Phi/\Psi$ [1, 18, 19, 145–147]. We follow the method of ref. [41], treating γ as a scale-independent factor, and consider a scenario in which Poisson’s equation for Φ is modified but Ψ is left unchanged. This amounts to making the substitution $\delta_\kappa = c_\kappa \delta_m$ in Eq. 3.10, where $c_\kappa = (1 + \gamma)/2$, since lensing probes the sum of the two potentials $\Phi + \Psi$. We approximate c_κ as a piecewise constant linear Eulerian bias, whose value in each redshift bin is treated as an independent parameter.




Figure 19. Error on the gravitational slip parameter γ from LSS×CMB lensing cross-correlations, LSS full shape data, and the primary CMB. For comparison, we show the constraint on γ from Planck [21] at $z \approx 1100$ (recombination) to the right of the plot.

In Fig. 19 we show constraints on γ_i , where γ_i is the value of the gravitational slip in the i^{th} redshift bin. In these constraints we include LSS×CMB lensing, full shape clustering data, the primary CMB, and marginalize over all parameters listed in Table 2. The direct constraints on γ_i come only from $C_\ell^{\kappa\kappa}$ and $C_\ell^{\kappa g_i}$, while the full shape and primary CMB data serve as priors on Λ CDM and the nuisance parameters. Thus the constraints in Fig. 19 should not be interpreted as independent measurements. Instead, they should be interpreted as a survey’s sensitivity to γ_i as measured by the full dataset (lensing cross-correlations and full shape clustering across the entire redshift range of the survey), assuming Λ CDM as the fiducial cosmology.

From Fig. 19 we see that spectroscopic surveys are capable of measuring $\gamma(z)$ to better than 10% out to redshift 6 with current CMB data. The improvement from upgrading to LiteBIRD + CMB-S4 from Planck + SO is $\sim 15\%$ for all LSS surveys at all redshifts of interest. We see that MegaMapper improves over MSE by 20 – 40% across $2 < z < 4$. These results are in rough agreement with ref. [41], who considered an LBG survey covering 1000 square degrees and found $0.02 < \sigma_\gamma < 0.20$ for $2 < z < 5$ after marginalizing over σ_8 and b_{LBG} . If we only marginalize over $\ln(A_s)$ and b_{LBG} , and reduce MegaMapper’s sky coverage to $f_{\text{sky}} = 0.024$, mimicking the approach of ref. [41], we find $0.09 < \sigma_\gamma < 0.23$ for $2 < z < 5$. The disagreement at low redshift is likely due to different binning schemes (we

use four redshift bins, while ref. [41] uses three), and differences in the assumed LBG number densities.

4.10 Extended base model and multi-extension fits

Our discussion so far focused on one parameter extensions of our base Λ CDM model. We now investigate how our constraints degrade as we extend our base model by a single parameter, or as we simultaneously vary multiple extensions.

We first consider the constraints on the Λ CDM parameters when the base model is extended by a single parameter. In Table 4 we show the percent increase to PSD + MegaMapper (using full shape data only, without lensing) constraints as we extend the base model to include M_ν , N_{eff} , Ω_k or f_{EDE} (with $\log_{10}(z_c) = 3.61$). In particular, we find that constraints on h , ω_c and ω_b worsen by roughly a factor of 2 when N_{eff} is allowed to vary. When the base model is extended to include M_ν or Ω_k , all Λ CDM constraints increase by less than 40%. Allowing f_{EDE} to vary, however, has a negligible (few percent) impact on Λ CDM constraints.

In Table 4 we also consider the impact of extending the base model on several one-parameter extensions. In particular, we find that $\{M_\nu, N_{\text{eff}}, \Omega_k\}$ are not internally degenerate – allowing any two of these three extensions to vary simultaneously worsens the constraints by less than 17%. We find a mild degeneracy between f_{EDE} and Ω_k ; varying the two simultaneously worsens the constraints on both parameters by $\approx 30\%$.

Extension	h	$\ln(A_s)$	n_s	ω_c	ω_b	τ	M_ν	N_{eff}	Ω_k	f_{EDE}
M_ν	39.7	29.5	2.7	26.6	0.0	35.3	–	11.5	12.5	16.2
N_{eff}	138.0	0.0	46.0	67.0	73.0	1.0	11.5	–	0.3	8.5
Ω_k	23.4	1.3	12.4	32.6	0.0	3.4	12.5	0.3	–	31.8
f_{EDE}	0.4	1.6	4.4	0.8	3.4	2.5	16.2	8.5	31.8	–

Table 4. Percent increase to PSD + MegaMapper constraints on Λ CDM (columns 2-7) and several one-parameter extensions (columns 8-11) when the base model is extended by a single parameter. The constraint on f_{EDE} is for fixed $\log_{10}(z_c) = 3.61$ and $\theta_i = 2.83$.

In Fig. 20 we consider changes in our constraints as we vary more than one extension simultaneously. The blue bars in the left panels of Fig. 20 shows the 1σ uncertainty on the Λ CDM parameters from PSD, which we compare to PSD + PUMA-32K, shown in orange (the analogous figure for MegaMapper looks nearly identical). Allowing for multiple extensions of the baseline model results in an overall degradation of the constraints, as shown by the different shading of the blue and orange bars. When the neutrino mass is marginalized over, this effect is the strongest on h , A_s and τ . This trend increases when simultaneously fitting for N_{eff} and Ω_k . The net improvement over PSD brought by future surveys is also reduced as more extensions are added to the base Λ CDM model, indicating that the set of observables we consider is not powerful enough to break the new degeneracies among parameters. Fitting to more than one extension to Λ CDM simultaneously without significant degradation to Λ CDM requires next generation space-based CMB surveys, such as LiteBIRD, or other measurements which are capable of constraining the primordial amplitude and optical depth to extremely high precision. These constraints are shown by the red and green bars respectively. As seen in Fig. 20, the constraints on τ and A_s from next generation CMB surveys are negligibly impacted when simultaneously fitting for multiple extensions. Thus no additional penalty is incurred to the remaining four Λ CDM parameters from the would-be worsened degeneracies with τ and A_s .




Figure 20. Constraints on Λ CDM and extensions from four different combinations of experiments, using primary CMB and clustering data only. We include a $\sigma(\tau) = 0.002$ prior for the green bars. *Left:* The shading/width denotes which parameters are allowed to vary when fitting for Λ CDM (e.g. for Base Λ CDM+ M_ν we keep N_{eff} and Ω_k fixed). *Right:* The shading denotes the number N of Parameter Extensions (NPE) that are simultaneously allowed to vary in our fits, starting with the relevant parameter on the x -axis and continuing along the cycle ($M_\nu, N_{\text{eff}}, \Omega_k$). E.g. for Ω_k , 2PE is equivalent to Ω_k, M_ν .

On the other hand, the constraints to our selected extensions to Λ CDM (right panel of Fig. 20) remain relatively unchanged while fitting to more than one extension. Just as before (Table 4), this suggests that $\{M_\nu, N_{\text{eff}}, \Omega_k\}$ are not internally degenerate.

4.11 Linear vs. non-linear modeling

Currently, the vast majority of forecasts assume linear theory and scale-independent bias (though see [63–65] for notable exceptions), and model the effects of FoG-like contributions with a specific functional form (typically a Gaussian or Lorentzian). That is, the typical linear model (LM) for the redshift-space galaxy power spectrum takes the form [148]:

$$P_{gg}^{\text{LM}}(\mathbf{k}) = \frac{(b + f\mu^2)^2}{1 + \sigma_v^2(k\mu)^2} P^{\text{lin}}(k) + N_0. \quad (4.5)$$

To explore the difference between this simplified linear model and our fiducial analysis, we create forecasts for DESI, MegaMapper and PUMA-32K (with optimistic foregrounds) following the procedure of §3.1 with Eq. (4.5) replacing the full model (FM), given by Eq. (3.3). A comparison between the constraints on Λ CDM and a selection of single parameter extensions is shown in Fig. 21. From Fig. 21 we see that the linear model generically predicts overly optimistic constraints: for both PSD and PSD + MegaMapper, the constraints using the full model are 20 – 40% larger than those using the linear model for most parameters, with n_s being a notable exception, where the linear and full models disagree by factors of 3 and 6 for the respective survey combinations. For PSD + PUMA-32K the disagreement is significantly worse: $\sigma(M_\nu)$ and $\sigma(N_{\text{eff}})$ from linear theory are overly optimistic by a factor of ~ 2 , while the constraints on n_s disagree by a factor of 12.

Leaving aside the fact that linear theory is clearly not sufficient for an unbiased estimate of cosmological parameters in the galaxy power spectrum, linear theory results in inaccurate forecasts for several reasons. Perhaps the most obvious is that the approximation of linear dynamics breaks down at intermediate to high k . Two beyond-linear effects are particularly important: the broadening of the BAO peak due to non-linear evolution and decorrelation




Figure 21. Constraints on Λ CDM and several one-parameter extensions when using the full model (FM), relative to the linear model (LM, red) or extended linear model (ELM, blue). In this plot we only include the primary CMB and full shape clustering data. Note that that largest radius in the two leftmost panels is 1.8, but 2.4 in the right panel. For the linear model, the values for n_s extend well beyond these limits: from left to right the ratios are 3.1, 6.8, and 11.6.

of the evolved field with the initial conditions [149–166]. In Fig. 7 we show a few logarithmic derivatives of the tracer power spectrum using both linear and non-linear theory at $z = 2.25$. The non-linear scale at this redshift is roughly $0.46 h \text{ Mpc}^{-1}$, and we see significant discrepancies (more than a few percent) by $k \sim 0.2 h \text{ Mpc}^{-1} < 0.5 k_{\text{nl}}$. In addition to linear theory breaking down on small scales, linear forecasting is inaccurate since the true angular structure is more complex than the $b + f\mu^2$ form allows, the bias model has too few degrees of freedom, linear forecasting ignores contributions from small scale physics (i.e. counterterms) and the FoG model is overly simplified.

The linear model can be trivially extended to address the latter two shortfalls: the effects of small-scale physics can be mimicked with the counterterm-like contribution $\mu^{2n} k^2 P^{\text{lin}}$, and the functional dependence of the FoG-like contribution can be relaxed via a Taylor expansion. Doing so results in the “extended linear model” (ELM):

$$P_{gg}^{\text{ELM}} = (b + f\mu^2)^2 P^{\text{lin}} + (\alpha_0 + \alpha_2 \mu^2 + \alpha_4 \mu^4) (k/k_*)^2 P^{\text{lin}} + N_0 + N_2 (\mu k)^2 + N_4 (\mu k)^4, \quad (4.6)$$

where we have suppressed the explicit \mathbf{k} dependence. A comparison of the ELM with the FM is shown in Fig. 21. For PSD (PSD + MegaMapper) we find a minor improvement in the ELM over the LM, with the mean disagreement¹² with the full model decreasing from 28% (28%) to 21% (25%). Most notable is the improved ELM prediction of $\sigma(n_s)$, which agrees with the full model to within 20% for both PSD and PSD + MegaMapper. For PSD + PUMA-32K, we find that the mean disagreement with the full model decreases from 44% to 25%, which is primarily driven by the significantly more accurate ELM predictions of $\sigma(M_\nu)$ and $\sigma(N_{\text{eff}})$, disagreeing by 20% and 40% with the full model, as opposed to a factor of 2. Likewise, the disagreement for $\sigma(n_s)$ decreases from a factor of 12 to 50%.

5 Observational challenges

In the previous section we neglected or delayed discussing the effects of several observational challenges that would arise in a more realistic treatment of future surveys. In this section we quantify some of these effects, summarizing the importance of the 21-cm foreground wedge,

¹²The “mean disagreement” $\equiv \sum_{\theta} [\sigma^{\text{FM}}(\theta) / \sigma^{\text{ELM}}(\theta) - 1] / 8$, where $\theta \in \{h, \ln(A_s), \omega_c, \omega_b, \tau, \Omega_k, N_{\text{eff}}, M_\nu\}$.

examining redshift uncertainties, the impact of a shallower u -band depth for LBG surveys, and the impact of a k_{\parallel} cut on constraints of primordial non-Gaussianity.

Throughout §4 we presented both optimistic and pessimistic forecasts for 21-cm surveys, representing the ability to mitigate foreground contamination in modes with low k_{\parallel} . Here we summarize the main results. Foreground assumptions can impact the precision of distance measurements perpendicular to the line of sight by more than a factor of 2, while measurements parallel to the line of sight are negligibly impacted. For PSD + PUMA-32K, the constraints on h , ω_c , A_{lin} and Ω_k vary by more than 50% depending on the foreground assumptions, while the constraints on N_{eff} and M_{ν} vary by more than 30%. For EDE the constraints can vary by more than a factor of 2. Finally, while the 21-cm constraints of $f_{\text{NL}}^{\text{Loc}}$ under optimistic foreground assumptions will not be competitive with a high- z galaxy survey, pessimistic foreground assumptions would inhibit a 21-cm measurement of $f_{\text{NL}}^{\text{Loc}}$ from scale-dependent bias altogether.




Figure 22. *Left:* Constraints on Λ CDM and several one-parameter extensions from MegaMapper full shape data vs redshift uncertainty, with $k_{\text{max}}/k_{\text{nl}} = 1$. The constraints are measured relative to $\sigma_z = 0$. *Right:* Change in constraints as a function of $k_{\text{max}}/k_{\text{nl}}$, relative to $k_{\text{max}}/k_{\text{nl}} = 1$. In both plots we include a prior on $\{\Lambda$ CDM, N_{eff} , Ω_k , $M_{\nu}\}$ from Planck + SO. The constraint on f_{EDE} is with fixed $\log_{10}(z_c) = 3.61$, while the constraint on A_{lin} is with fixed $\omega_{\text{lin}} = 147 h^{-1} \text{Mpc}$.

As previously discussed in §3.7, at high redshifts it becomes increasingly difficult to obtain small redshift errors, while the effects of redshift errors become increasingly more important due to the non-linear scale shifting to higher k . In Fig. 22 we show the impact of redshift errors on the constraints of Λ CDM and several one-parameter extensions from MegaMapper full shape data. We find that redshift errors negligibly impact our results for $\sigma_z/(1+z) < 10^{-3}$, only degrading the constraints at the few percent level. However, when $\sigma_z/(1+z)$ nears 0.3% we begin to see non-negligible ($\sim 25\%$) effects on some parameters, most notably M_{ν} , f_{EDE} and A_{lin} . As shown in the right panel of Fig. 22, including redshift errors is roughly equivalent to enforcing a strict $k_{\text{max}}/k_{\text{nl}}$ cut, with $\sigma_z/(1+z) = 0.01$ (0.003) corresponding to roughly $k_{\text{max}}/k_{\text{nl}} = 0.3$ (0.6) for our fiducial MegaMapper survey.

To estimate the impact of a lower u -band depth on MegaMapper constraints, we consider the idealized MegaMapper-like survey described in §2, which we split into four redshift bins spanning $2 < z < 5$. To mimic the effect of a lower u -band depth, we vary m_{UV} in the two central redshift bins (spanning $2.75 < z < 4.25$) between 23 and 23.5. We leave $m_{\text{UV}} = 24.5$ in the lowest and highest redshift bins unchanged, and assume a mean 50% target efficiency ($\bar{n}/n_{\text{LBG}} = 0.5$) in all redshift bins, regardless of m_{UV} . As shown in the right panel in Fig. 23,

the constraints from Planck + SO + MegaMapper on Λ CDM and common one parameter extensions (N_{eff} , M_ν , Ω_k) increase by less than $\approx 15\%$ as the u -band depth is decreased from 24.5 to 23, while the constraints on f_{EDE} (at fixed $\log_{10}(z_c) = 3.61$) and A_{lin} (at fixed $\omega_{\text{lin}} = 147 h^{-1} \text{Mpc}$) increase by less than 25%.




Figure 23. *Left:* Relative improvement to Planck + SO + MegaMapper-like parameters as (MegaMapper) area is traded for depth, for fixed observing time. We assume that $m_{\text{UV}}(t) = 24.5 + 2.5 \log_{10}(\sqrt{t})$ and $f_{\text{sky}}(t) = 0.34/t$, and scan through $0.4 < t < 8$. *Right:* Degradation to Planck + SO + MegaMapper-like constraints with a shallower u -band depth. Here the MegaMapper-like survey has a default 24.5 limiting magnitude, and detects 50% of the LBGs for a given m_{UV} .

We also perform a simplistic area vs depth trade-off study for the same MegaMapper-like survey, which we discuss here. For fixed observing time, we assume the sky coverage and the change to the limiting apparent magnitude scale as $f_{\text{sky}} \propto t^{-1}$ and $\delta m_{\text{UV}} \propto 2.5 \log_{10}(\sqrt{t})$ respectively, which should be a very good approximation given that MegaMapper will be deep into the sky-noise limit. The fiducial survey corresponds to $t = 1$. In principle, changing the limiting magnitude changes the bias of the LBG sample. We account for this effect by rescaling $n(m = 24.5, z) \rightarrow n(m, z)b^2(m, z)/b^2(m = 24.5, z)$ in our covariance.




Figure 24. Constraints on the amplitude of local primordial non-Gaussianity with a cut on k_{\parallel} modes. In the left panel we set $k_{\text{max}} = 0.02 h \text{Mpc}^{-1}$, while in the right panel we set $k_{\text{max}} = 0.05 h \text{Mpc}^{-1}$. In both panels we take $k_{\text{min}} = 0.001 h \text{Mpc}^{-1}$.

As shown in the left panel of Fig. 23, the fiducial MegaMapper sky coverage ($f_{\text{sky}} = 0.34$) is well optimised for the parameters considered in this work. In particular, we find that the

fiducial MegaMapper sky coverage yields constraints that are within 5% of their optimal values.

Finally we consider the impact of a k_{\parallel} cut, which could be imposed to remove systematic errors, on constraints of local primordial non-Gaussianity. For our constraints we follow the method described in §4.6 with $k_{\min} = 0.001 h \text{ Mpc}^{-1}$, and remove modes whose component along the line of sight is smaller than some k_{\parallel} cut, which we vary between 0 and $0.01 h \text{ Mpc}^{-1}$. Our results are illustrated in Fig. 24. With a conservative $k_{\max} = 0.02 h \text{ Mpc}^{-1}$, we find that MegaMapper constraints degrade from $\sigma(f_{\text{NL}}^{\text{Loc}}) = 0.7$ to $\sigma(f_{\text{NL}}^{\text{Loc}}) = 10$ as the k_{\parallel} cut is increased from 0 to $0.01 h \text{ Mpc}^{-1}$. For a more optimistic $k_{\max} = 0.05 h \text{ Mpc}^{-1}$ these constraints degrade from 0.7 to 4.

6 Conclusions

Advances in experimental techniques make it possible to map the high redshift Universe at high fidelity in the near future. This will increase the observed volume by many-fold, while providing unprecedented access to very large scales, allowing the reconstruction of an order of magnitude more “primordial modes”. Such surveys would improve measurements of the expansion history and spatial curvature of the Universe, primordial non-Gaussianity, features in the power spectrum, models of inflation, dark energy and modified gravity, light relics and neutrino masses. Recently developed theoretical models, well suited to the quasi-linear clustering of the early Universe, provide a firm basis for forecasting the science reach of such surveys including such complexities are non-linear evolution, redshift-space distortions and non-linear bias. In this paper we have presented such forecasts for a handful of proposed experiments (§2) that illustrate the potential of next-generation large-scale structure surveys to inform our theories of fundamental physics and discussed some of the limitations and synergies with other probes. We used a Fisher matrix formalism (§3) to forecast errors on cosmological parameters from the experiments individually and in combination with other probes, such as primary CMB anisotropies and CMB lensing. Baryon acoustic oscillations (BAO) are treated as a special case, as are methods for determining the amplitude of clustering using a fixed shape of the power spectrum (e.g. refs. [105, 106]).

Section 4.1 shows that such surveys could map the expansion history and spatial curvature of the Universe, in a relatively model independent fashion and with high precision, over more than 90% of cosmic time ($z < 6$). Fig. 10 shows that MegaMapper could provide sub-percent measures of the angular diameter distance to $z \approx 5$ while a PUMA-like experiment could reduce the errors by a further factor of 2 under optimistic assumptions for foregrounds. The line-of-sight distance, i.e. the expansion rate, can be determined with similar accuracy with a PUMA-like experiment providing $< 0.5\%$ precision to $z \approx 6$ for a wide range of foreground assumptions.

Section 4.3 shows that future high redshift surveys will enable percent-level measurements of structure growth out to redshift 5. Fig. 11 shows that PUMA could achieve 5 – 10% measurements of $\sigma_8(z)$ across $2 < z < 6$, while MegaMapper could reduce these errors to less than 1% across $2 < z < 5$. 21-cm surveys are severely limited by a degeneracy with the mean 21-cm brightness temperature, while high redshift spectroscopic surveys benefit greatly from the addition of CMB lensing cross-correlations, which decrease the error on $\sigma_8(z)$ by a factor of 2.

The ΛCDM model with minimal-mass neutrinos provides a highly predictive and rigid model that can explain the majority of cosmological observations [21, 106] and thus makes

a natural target for forecasts (§4.2). Table 3 and Figs. 9 and 20 present a summary of our findings. Upcoming surveys like DESI and Euclid, in combination with Planck, will improve current measurements of ω_c and h by a factor of two, but will largely exhaust the information available from the low redshift power spectrum. Adding high- z redshift measurements allows further improvements by more than a factor of 3, exceeding the reach of even ambitious next-generation CMB experiments. The improvement in the determinant of the Fisher matrix for base Λ CDM in going from PSD to PSD + MegaMapper (PSD + PUMA-32K) is a factor of 2×10^4 (7×10^4).

Massive neutrinos provide one of the best motivated extensions of the minimal Λ CDM model (§4.4), and large-scale structure provides the tightest constraints on the sum of the neutrino masses [114]. Table 3 shows that constraints on M_ν tighten significantly with the inclusion of high-redshift LSS data. While the forecasts depend sensitively on the smallest scales that can be used (§4.4; Fig. 12) the next-generation of CMB surveys coupled with a next-generation, high-redshift LSS survey should comfortably achieve $\sigma(M_\nu) = 20$ meV. This could be further improved by a better measurement of τ (from the CMB or other means) and the inclusion of CMB lensing, which aids in degeneracy breaking (see Fig. 27).

Light relics (§4.5) suppress structure on small scales and thus can also be well constrained by high-redshift large-scale structure surveys (Table 3 and Fig. 13). We find that both MegaMapper and PUMA are nearly cosmic variance limited for this measurement and capable of achieving $\sigma(N_{\text{eff}}) \simeq 0.02$ in combination with future CMB experiments (depending upon the k_{max} assumed: Table 3 and Fig. 13). Such a measurement could detect light vectors at 3σ , regardless of their freeze-out temperature, and would be capable of detecting (or ruling out) all light relics with a freeze-out temperature smaller than the QCD phase transition (~ 0.2 GeV) at 5σ significance. Interestingly we find that a proper marginalization over nuisance parameters (such as non-linear biases and counterterms) is important for forecasting $\sigma(N_{\text{eff}})$, with constraints underestimated by $\simeq 40\%$ when these parameters are held fixed (as is often done).

One window into the very early Universe and the physics of inflation is the study of primordial non-Gaussianity (§4.6). The Planck constraints $f_{\text{NL}}^{\text{loc}} = -0.9 \pm 5.1$ are within a factor of 3 of the cosmic variance limit, so pushing to $\sigma(f_{\text{NL}}^{\text{loc}}) \lesssim 1$ requires LSS data. We find a high- z LSS survey would be capable of a 5-fold improvement over a cosmic variance limited CMB survey, with $\sigma(f_{\text{NL}}^{\text{loc}}) \sim 0.8$, just from the power spectrum alone. Further improvements could be possible by including information from the bispectrum or from including CMB lensing.

A second window into the physics of inflation from large-scale structure surveys is to search for primordial features (§4.7) in the power spectrum. The space of possible features is large, so as an example we show that a future LSS survey could constrain the amplitude of a sinusoidal modulation of $P(k)$ to better than 0.1% over a broad range of frequencies (Fig. 15). A precise measurement of the power spectrum shape can also be used to place constraints on (or detect) short-term deviations from matter or radiation domination, since such periods will change the shape of the power spectrum due to differential growth of modes. Thus the shape of the matter power spectrum has the potential to detect sub-percent fractional amounts of Early Dark Energy to $z \sim 10^5$ (§4.8), probing Dark Energy all the way to when the Universe was only a few years old (Fig. 18).

The precision of these measurements, combined with CMB observations, also has the promise of tightening constraints on the gravitational slip, which is equal to one in General Relativity but generically departs from unity in other models. For the simplest case of a

scale-independent slip, we show that future high redshift experiments could constrain γ to a few percent (Fig. 19) with the tightest constraints coming from MegaMapper near $z \approx 2 - 3$.

We find that inclusion of models that go beyond constant bias and linear theory is important for accurately assessing the potential of future experiments to constrain parameters. Neglecting the complexities of bias or fixing too many nuisance parameters can artificially inflate the precision of parameter measurements by a factor of up to 12 in Λ CDM models (§4.11 and Fig. 21).

At higher redshift the non-linear scale shifts to higher k , and the impact of fingers of god on the redshift-space clustering also declines. To make use of the extra information available along the line of sight requires an increase in redshift precision (§3.7; Figs. 5 and 8). Achieving this precision observationally may be challenging if broad lines, or lines with significant radiative transfer effects, are used. In §5 we discuss the impact of redshift uncertainties on a MegaMapper-like survey, and find that our constraints degrade by less than 15% for $\sigma_z/(1+z) < 0.002$, which corresponds to a systematic velocity of 600 km s^{-1} .

Studies of linear or quasi-linear large-scale structure with redshift surveys and the CMB currently provide our tightest constraints on cosmology and fundamental physics. Pushing the redshift and volume frontier will provide guaranteed, significant improvements in the state-of-the-art in a manner that can be reliably forecast and optimized. This paper presents a first step in this direction, considering a number of proposed experiments that would map the Universe in the epochs before cosmic noon.

Acknowledgments

We would like to thank Arjun Dey and David Schlegel for helpful discussions on redshift errors for high redshift galaxies and Kyle Dawson for conversations and calculations regarding SpecTel. We also thank Shi-Fan Chen, Mikhail Ivanov, Marko Simonovic, Juna Kollmeier, Eric Linder, Emmanuel Schaan, Uroš Seljak, Leonardo Senatore, Anže Slosar and Michael Wilson for useful discussions during the preparation of this manuscript. N.S. is supported by the NSF. S.F. is supported by the Physics Division of Lawrence Berkeley National Laboratory. M.W. is supported by the DOE and the NSF. This research has made use of NASA’s Astrophysics Data System and the arXiv preprint server. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

A Primordial figure of merit (N_{modes})

Here we describe a new figure of merit for “counting reconstructable linear modes.” This is similar to the effective volume, but integrated to a k_{max} that incorporates the decorrelation of the observed nonlinear field with the initial conditions at high k , and hence loss of primordial information on small scales. This loss of information can be modeled as a Gaussian “propagator” $G(\mathbf{k})$. That is, we split the observed nonlinear overdensities $\delta_F(\mathbf{k}) = G(\mathbf{k})\delta_L(\mathbf{k}) + d(\mathbf{k})$ into two pieces: one which is correlated with the linear overdensities, $\delta_L(\mathbf{k})$, and one which isn’t $\langle \delta_L(\mathbf{k})d(\mathbf{k}') \rangle = 0$. The Gaussian propagator is given by:

$$G(k, \mu) \equiv \frac{\langle \delta_F \delta_L \rangle}{\langle \delta_L \delta_L \rangle} \simeq \exp \left[-\frac{1}{2} \left(k_{\perp}^2 + k_{\parallel}^2 \{1 + f\}^2 \right) \Sigma^2 \right], \quad (\text{A.1})$$

where Σ is set by the rms displacement within the Zel’dovich approximation, as in Eq. (3.8).

In the sample variance limit, the relative uncertainty of the linear power spectrum's amplitude A is given by $\sqrt{2/N_{\text{modes}}}$. We can calculate this uncertainty from the Fisher matrix:

$$F_{ii} = \frac{f_{\text{sky}}}{2} \int_{z_{\text{min}}}^{z_{\text{max}}} dz \frac{dV}{dz} \int_{k-\text{wedge}(z)}^{\infty} \frac{d^3\mathbf{k}}{(2\pi)^3} \left(\frac{\partial_{\theta_i} P_F(\mathbf{k}, z)}{P_F(\mathbf{k}, z) + N(z)} \right)^2, \quad (\text{A.2})$$

where $k-\text{wedge}(z)$ defines the lower limit of integration. The uncertainty on the amplitude is $\delta A/A = \delta P_L(\mathbf{k})/P_L(\mathbf{k}) = 1/\sqrt{F_{AA}}$, so that $N_{\text{modes}} = 2F_{AA}$. This defines the figure of merit N_{modes} .

Let $\hat{A} = P_L(\mathbf{k})/\hat{P}_L(\mathbf{k})$ be an estimator for the amplitude of the linear power spectrum $P_L(\mathbf{k}) = (b + f\mu^2)^2 P_m(k)$, where $\hat{P}_L(\mathbf{k})$ is any (unbiased) estimator for the linear galaxy power spectrum in redshift space, and $P_m(k)$ is the linear matter power spectrum. Since $P_F(\mathbf{k}) = G^2(\mathbf{k})P_L(\mathbf{k}) + P_d(\mathbf{k})$ does not explicitly depend on $\hat{P}_L(\mathbf{k})$, the derivative of the non-linear power spectrum with respect to \hat{A} is simply¹³ $G^2(\mathbf{k})\hat{P}_L(\mathbf{k})$. This derivative should be evaluated at the fiducial cosmology, so $\partial_A P_F(\mathbf{k}) = G^2(\mathbf{k})P_L(\mathbf{k})$. Putting this together gives:

$$N_{\text{modes}} = f_{\text{sky}} \int_{z_{\text{min}}}^{z_{\text{max}}} dz \frac{dV}{dz} \int_{k-\text{wedge}(z)}^{\infty} \frac{d^3\mathbf{k}}{(2\pi)^3} \left(\frac{G^2(\mathbf{k}, z)P_L(\mathbf{k}, z)}{P_F(\mathbf{k}, z) + N(z)} \right)^2. \quad (\text{A.3})$$

The calculation of the CMB equivalent of N_{modes} is almost identical. We include the TT, TE, and EE power spectra $\boldsymbol{\mu}_\ell = (C_\ell^{\text{TT}}, C_\ell^{\text{TE}}, C_\ell^{\text{EE}})^T$ in our Fisher forecast, which has covariance matrix:

$$\mathbf{C}_{\ell, \ell'} = \frac{2\delta_{\ell, \ell'}^K}{(2\ell + 1)f_{\text{sky}}} \begin{pmatrix} (C_\ell^{\text{TT}} + N_\ell)^2 & (C_\ell^{\text{TT}} + N_\ell) C_\ell^{\text{TE}} & (C_\ell^{\text{TE}})^2 \\ (C_\ell^{\text{TT}} + N_\ell) C_\ell^{\text{TE}} & \frac{1}{2} [(C_\ell^{\text{TE}})^2 + (C_\ell^{\text{TT}} + N_\ell)(C_\ell^{\text{EE}} + 2N_\ell)] & (C_\ell^{\text{EE}} + 2N_\ell) C_\ell^{\text{TE}} \\ (C_\ell^{\text{TE}})^2 & (C_\ell^{\text{EE}} + 2N_\ell) C_\ell^{\text{TE}} & (C_\ell^{\text{EE}} + 2N_\ell)^2 \end{pmatrix}, \quad (\text{A.4})$$

where

$$N_\ell = \Delta_T^2 \exp\left(\frac{\theta_{\text{FWHM}}^2 \ell^2}{8 \ln 2}\right) \quad (\text{A.5})$$

accounts for the instrumental noise. Here Δ_T is the noise level of the experiment and θ_{FWHM} is the full-width at half-maximum (FWHM) of the beam in radians. Given $\boldsymbol{\mu}_\ell$ and $\mathbf{C}_{\ell, \ell'}$, the Fisher matrix is simply:

$$F_{ij} = \frac{1}{2} \sum_{\ell, \ell_1, \ell_2, \ell_3} \text{Tr} \left[\mathbf{C}_{\ell, \ell_1}^{-1} \frac{\partial \mathbf{C}_{\ell_1, \ell_2}}{\partial p_i} \mathbf{C}_{\ell_2, \ell_3}^{-1} \frac{\partial \mathbf{C}_{\ell_3, \ell}}{\partial p_j} \right] + \sum_{\ell, \ell'} \frac{\partial \boldsymbol{\mu}_\ell^T}{\partial p_i} \mathbf{C}_{\ell, \ell'}^{-1} \frac{\partial \boldsymbol{\mu}_{\ell'}}{\partial p_j}. \quad (\text{A.6})$$

The second term dominates since the summand is weighted by the number of modes, whereas in the first term all factors of $2\ell + 1$ cancel. Just as before, we take $N_{\text{modes}} = 2F_{AA}$:

$$N_{\text{modes}} \simeq 2 \sum_{\ell, \ell'} \boldsymbol{\mu}_\ell^T \mathbf{C}_{\ell, \ell'}^{-1} \boldsymbol{\mu}_{\ell'}. \quad (\text{A.7})$$

Due to the bias from foreground contamination, C_ℓ^{TT} will realistically be measurable out to $\ell_{\text{max}} \simeq 3000$. We enforce this in our calculation of N_{modes} by multiplying C_ℓ^{TT} by 10^6 in our covariance matrix when $\ell > 3000$.

¹³This is not true in a strict mathematical sense. In reality the decorrelated piece $P_d(\mathbf{k})$ is still dependent on the initial conditions, although the manner in which it depends on $P_L(\mathbf{k})$ is highly non-trivial and subdominant, so we neglect it in our derivative.

Our definition of N_{modes} can be easily generalized to represent a survey’s ability to measure the amplitude of the linear matter power spectrum (at $z = 0$) in a specific k -bin. Following the same argument as before, let $\hat{A}(k) = P_m(k)/\hat{P}_m(k)$ be an estimator for the amplitude, where $P_m(k)$ is the linear matter power spectrum at $z = 0$. Then $\partial_{\hat{A}(k')} P_F(\mathbf{k}, z) = G^2(\mathbf{k}, z) T^2(k, z) (b + f\mu^2)^2 \hat{P}_m(k) \Theta(\Delta k - |k - k'|)$, where $T(k, z)$ is the transfer function, $\Theta(x)$ is the Heaviside step function, and Δk is the width of the k -bin. When evaluated at the fiducial cosmology, this yields $\partial_{\hat{A}(k')} P_F(\mathbf{k}, z) = G^2(\mathbf{k}, z) P_L(\mathbf{k}, z) \Theta(\Delta k - |k - k'|)$. Just as before, we interpret $2F_{A(k)A(k)}$ as the effective number of sample variance limited modes that measure the amplitude:

$$N_{\text{modes}}(k) = f_{\text{sky}} \int_{z_{\text{min}}}^{z_{\text{max}}} dz \frac{dV}{dz} \frac{k^2 \Delta k}{2\pi^2} \left(\frac{G^2(\mathbf{k}, z) P_L(\mathbf{k}, z)}{P_F(\mathbf{k}, z) + N(z)} \right)^2. \quad (\text{A.8})$$

We make use this expression in Fig. 3, where the error in each bin is $P_m(k) \sqrt{2/N_{\text{modes}}(k)}$.

One note on the interpretation of N_{modes} is in order here: we have defined N_{modes} to be the number of modes recoverable that are correlated with the initial conditions. Therefore it is a figure of merit for how well we can constrain the primordial potential or other statistical properties of the initial conditions. N_{modes} as defined here does not take into account the sensitivity of a given observable to the parameters of interest, so it should not be taken as the only comparison metric, especially between high and low-redshift probes (such as CMB vs LSS). For example, it is very possible for an LSS experiment with lower N_{modes} than the CMB to be a lot more sensitive to Dark Energy or neutrino masses, since it performs measurements on a range of redshifts where they are dynamically important.

For this reason, while a direct comparison between CMB and LSS experiments based on N_{modes} alone might be misleading, we believe that it provides a useful metric to compare different survey designs covering a similar redshift range. Since it represents a proxy for “primordial information” available, survey optimization aimed at maximizing it can ensure better sensitivity to the physics of the early Universe.

B Covariance approximations

In the main text we have followed the “usual” approach of approximating the covariance matrix by the Gaussian (disconnected) piece. In Eq. (3.6) we also assume that the survey has a uniform number density in each redshift bin. As emphasized by ref. [83], a more realistic approach would include the impact of window functions, galaxy weights, etc. Inclusion of those effects leads to a covariance with the same functional form

$$C \propto (P^2 + 2P\bar{n}^{-1} + \bar{n}^{-2})/V \quad (\text{B.1})$$

with the substitutions

$$V \rightarrow I_{22}^2/I_{44}, \quad \bar{n}^{-1} \rightarrow I_{34}/I_{44}, \quad \bar{n}^{-2} \rightarrow I_{24}/I_{44}, \quad (\text{B.2})$$

where the I_{ij} ’s are integrals over the number density [83]:

$$I_{ij} = \int d^3x \bar{n}^i(\mathbf{x}) w^j(\mathbf{x}) \quad \text{with} \quad w(\mathbf{x}) = \frac{1}{1 + \bar{n}(\mathbf{x}) P_{\text{fid}}} \quad (\text{B.3})$$

The 2D model is only slightly more complex, with different damping across and along the line of sight and redshift-space distortions included assuming linear theory. To include the effects of reconstruction the damping parameters (Σ) are reduced by ≈ 0.5 . By contrast our method uses the Zeldovich-based model of ref. [67]. This model automatically accounts for the differences in damping of the different components, more complex bias, non-linear redshift-space distortions and the manner in which reconstruction reduces the damping. By using a forward model of the reconstructed power spectrum, our forecast is of the standard Fisher form wherein a parameter error is forecast through the theory derivative rather than requiring any modification of the procedure.

However, our method is actually quite close (within $\sim 15\%$) to the standard approach in practice, as shown in Fig. 26. The damping of the BAO signal by non-linear evolution is very close to Gaussian, as assumed in the standard approach, and while reconstruction affects each of the damping terms differently along and across the line of sight the major impact is to reduce the scale of the damping by approximately a factor of 2 (indeed it was a perturbative analysis of BAO reconstruction that suggested this approximation).




Figure 26. In the top two panels we compare our distance constraints (solid), calculated using the method of §3.6, to those found using the standard procedure of ref. [107] (dashed). The ratio of the two methods is shown in the bottom two panels.

D Lensing auto- vs. cross-correlation

Here we comment on the information content of the CMB lensing auto power spectrum, vs the cross-correlation with galaxies. Given a spectroscopic survey with sufficiently high number densities that covers the majority of the redshift range where the CMB lensing kernel is large, one can in principle reconstruct $C_\ell^{\kappa\kappa}$ from the galaxy auto-correlation and galaxy-convergence cross-correlation. This is easily seen in linear theory for a noiseless galaxy survey, so that the number density $n \rightarrow \infty$. In this limit we can make the redshift bins arbitrarily small, and approximate the integral expression for $C_\ell^{\kappa g_i}$ and $C_\ell^{g_i g_i}$ (Eq. (3.10)) at a single comoving distance χ_i . Under these approximations, $C_\ell^{\kappa\kappa}$ is related to $C_\ell^{\kappa g_i}$ and $C_\ell^{g_i g_i}$ via

$$C_\ell^{\kappa\kappa} = \sum_i (C_\ell^{\kappa g_i})^2 / C_\ell^{g_i g_i}, \quad (\text{D.1})$$

where i runs over redshift bins covering $0 < z < z_*$. Thus $C_\ell^{\kappa g_i}$, $C_\ell^{g_i g_i}$ carries at least the same information as $C_\ell^{\kappa\kappa}$ in this idealized case. Since $C_\ell^{\kappa g_i}$, $C_\ell^{g_i g_i}$ can in principle be measured to a higher signal-to-noise than $C_\ell^{\kappa\kappa}$, and can be used to measure time evolution, one might naively

expect that parameter constraints are driven by $C_\ell^{\kappa g_i}$, $C_\ell^{g_i g_i}$. However, in our forecasts we find this not to be the case, as illustrated in Fig. 27. The contours from an SO measurement of the convergence auto-correlation (red) are significantly tighter than the contours from galaxy auto- and cross-correlations, as measured by SO and our fiducial PUMA-32K survey (light blue).




Figure 27. Constraints from cross-correlations of SO lensing with HIRAX-ideal (grey), PUMA-ideal (light blue), a version of PUMA-ideal that extends down to $z = 0.2$ (dark blue), and PUMA-ideal with non-linear nuisance terms held fixed (green). In red we plot the constraints from the SO lensing convergence autospectrum. To illustrate how the degeneracies from lensing and full shape measurements differ, we plot the constraints from PUMA-32K full shape data in yellow. All contours share the same Planck + SO prior on Λ CDM+ M_ν .

There are multiple ways in which this simplistic argument goes wrong: (1) the assumption that $C_\ell^{\kappa g_i}$, $C_\ell^{g_i g_i}$ can be measured to a high signal-to-noise ratio depends on the details of the survey, (2) a redshift survey might not cover the entire region where the lensing kernel is large, and (3) forecasts that include galaxy auto- and cross-correlations marginalize over

Parameter	b_2	b_s	α_0	α_2	α_4	N_2	N_4
Default value	0.93	-0.70	-8.65	0	0	-728	0
New value	1.8	-1.4	0	5	5	0	318

Table 5. Default and new fiducial values for the nuisance parameters of the MiniMapper survey.

galaxy biases, counterterms, and a Poisson-like stochastic contribution; however, including $C_\ell^{\kappa\kappa}$ in a forecast introduces no new nuisance parameters. This is by no means an exhaustive list, since Eq. (D.1) is no longer true for a survey with finite noise, redshift bins of finite extent, when magnification bias is included, or when modeling the tracer beyond linear theory, etc.

To test which of these effects causes the largest degradation to constraints from $C_\ell^{\kappa g_i}$ and $C_\ell^{g_i g_i}$, we consider idealized versions of the HIRAX and PUMA-32K surveys which have no foreground wedge and are assumed to measure T_b to infinite precision (so that it is held fixed), dubbed HIRAX-ideal and PUMA-ideal, enabling cross-correlations with CMB lensing. We take HIRAX-ideal and PUMA-ideal as representative (toy-model) low noise surveys at high redshift. As illustrated in Fig. 27, we see a negligible improvement to parameter constraints when upgrading from HIRAX-ideal (grey) to PUMA-ideal (light blue). Thus (1) is not the problem, and PUMA-ideal has sufficiently low noise levels for high fidelity measurements of $C_\ell^{\kappa g_i}$ and $C_\ell^{g_i g_i}$. We also see that extending PUMA-ideal to $z = 0.2$ (light vs. dark blue) has a small impact on the constraints, suggesting that PUMA’s fiducial redshift range ($2 < z < 6$) sufficiently covers the region where the CMB lensing kernel is large, eliminating (2). Instead, the primary culprit is (3). From Fig. 27 we see that the covariance from the convergence auto-correlation (red) is nearly identical to the covariance from the galaxy auto- and cross-correlation (green) when the nuisance terms are fixed.

We should stress that the value of cross-correlations extends considerably beyond constraining standard cosmological parameters. While $C_\ell^{\kappa\kappa}$ depends on the integrated matter fluctuations over a wide redshift range, as discussed in Sec. 4.3, they allow for a tomographic reconstruction of the amplitude of structure $\sigma_8(z)$ that only depends on the local amplitude. This can provide important model-independent tests beyond the rigidity of Λ CDM, considerably opening up the discovery space.

E Empirical checks

As discussed in §3.1, we only integrate over modes where the N_2 term ($N_0\sigma_v^2(k\mu)^2$) is less than 20% of the fiducial power. This “ N_2 cut” is only relevant for MSE and MegaMapper, whose number densities are sufficiently low enough for the N_2 term to contribute a non-negligible amount to the fiducial power for $k < k_{\text{nl}}$. As we show in Fig. 28, our results are largely insensitive to the choice of N_2 cutoff. For any cutoff above 4% the constraints from Planck + SO + MegaMapper change by at most 10%.

To quantify the importance of nuisance parameters’ fiducial values, we consider a “MiniMapper” survey that observes LBGs in a single redshift bin covering $2 < z < 3$ with $f_{\text{sky}} = 0.34$. We assume that MiniMapper has 100% efficiency for a limiting magnitude $m_{\text{UV}} = 24.5$. The default values of MiniMapper’s nuisance terms (following the conventions in §3.1) are listed in the first row of Table 5. The constraints from Planck + MiniMapper on $\{h, \omega_c, \tau, \Omega_k, N_{\text{eff}}, M_\nu\}$ with these default values are shown in the last row of Table 6. The remaining rows of Table 6 show the constraints when the fiducial value of each of the seven




Figure 28. Planck + SO + MegaMapper (full shape only) constraints as the N_2 cutoff is varied. This is with $\log_{10}(z_c) = 3.61$ and $\omega_{\text{lin}} = 147 h^{-1} \text{Mpc}$.

Parameter changed	$10^4 h$	$10^4 \omega_c$	$10^3 \tau$	$10^3 \Omega_k$	$10^3 N_{\text{eff}}$	M_ν [meV]
b_2	25.1	5.82	6.12	0.71	72.35	34.02
b_s	22.77	5.42	4.96	0.69	77.17	28.23
α_0	19.58	4.79	5.35	0.69	81.02	31.96
α_2	19.4	4.75	5.34	0.69	80.38	31.85
α_4	19.23	4.71	5.34	0.69	79.88	31.83
N_2	19.75	4.83	5.32	0.69	81.35	31.73
N_4	19.69	4.82	5.31	0.69	81.23	31.7
Default	19.68	4.82	5.31	0.69	81.23	31.7

Table 6. Constraints on selected Λ CDM parameters and extensions for Planck + MiniMapper. In each row we change the relevant nuisance parameter’s fiducial value to the new value listed in Table 5. The fiducial values of the remaining nuisance parameters are left unchanged. In the last row we show the constraints when all of the nuisance terms fiducial values’ are set to the default values in Table 5.

nuisance terms is individually varied, with the new fiducial values listed in the second row of Table 5. We find that varying the values of $\{\alpha_0, \alpha_2, N_2, \alpha_4\}$ has a negligible ($< 6\%$) impact on our results, while varying the fiducial values of $\{b_2, b_s\}$ has less than a 27% impact. We should be able to refine our fiducial values as we learn more about the clustering of these high redshift samples.

We have also checked the sensitivity of Early Dark Energy constraints to the fiducial value of f_{EDE} . We find that the constraint on f_{EDE} (for $\log_{10}(z_c) = 3.5$ and $\theta_i = 2.83$) from Planck + MiniMapper changes by less than 3% when the fiducial value of f_{EDE} is increased from 0 to 0.02.

F Python implementation

The Python code used to create these forecasts (FishLSS) is publicly available¹⁵. We include an example Jupyter notebook that walks the user through using the code in detail, and provide a brief overview of the organizational structure below.

¹⁵<https://github.com/NoahSailer/FishLSS>

The primary module is `fisherForecast.py`, which takes a cosmology (an instance of `CLASS`) and experiment (an instance of `experiment.py`) as inputs, and creates a Fisher forecast for a user-specified set of parameters and observables. `fisherForecast.py` contains functions for calculating derivatives of power spectra, computing the 21-cm foreground wedge (under optimistic or pessimistic assumptions), computing Fisher matrices for both 3D and 2D clustering measurements, combining the Fisher matrices from two uncorrelated measurements, and calculating both N_{modes} and $N_{\text{modes}}(k)$. `fisherForecast.py` has three helper modules – `experiment.py`, `twoPoint.py`, and `twoPointNoise.py` – which we describe below.

`experiment.py` specifies the experimental configuration of the survey, including the redshift range, redshift bins, linear bias, nonlinear biases (optional), redshift errors, velocity dispersion for FoG-like contributions, and sky coverage. For 21-cm surveys, additional details of the interferometric array are stored in `experiment.py`, such as the number of detectors, fill factor and observing time. We include a set of flags (`LBG`, `HI`, `HalpHa`, `ELG`, `Euclid`, `MSE`) which may be switched to `True` should the user wish to use the default biases and number densities of a specific survey described in §2.

Power spectra are computed with `twoPoint.py`. In particular, `twoPoint.py` has several functions for computing the fiducial linear biases for the surveys and samples considered in this work, as well as functions for computing both the matter and galaxy power spectra in redshift space ($P_{mm}(\mathbf{k})$, $P_{gg}(\mathbf{k})$); the real space galaxy-matter cross spectrum ($P_{gm}(k)$); the CMB lensing auto-spectrum, cross-correlation with galaxies, and the projected galaxy auto-spectrum ($C_{\ell}^{\kappa\kappa}$, $C_{\ell}^{\kappa g}$, C_{ℓ}^{gg}); and the reconstructed power spectrum. These functions require the installation of both `CLASS`¹⁶ [68] and `velocileptors`¹⁷ [41], which in turn requires `pyFFTW`¹⁸.

We calculate covariances (at the \mathbf{k} or ℓ -level) with `twoPointNoise.py`. In particular, `twoPointNoise.py` contains functions for computing the Poisson shot noise of each of the samples considered in this work, as well as functions for computing Eqs. (3.6) and (3.14).

While we do not include our code for computing the CMB priors, the Fisher matrices used for Planck, SO, S4, and LiteBIRD can be found in the `input` folder.

References

- [1] A. Slosar, R. Mandelbaum, and D. Eisenstein, *Dark Energy and Modified Gravity*, *Bull. Am. Astron. Soc.* **51** (May, 2019) 97, [[arXiv:1903.12016](https://arxiv.org/abs/1903.12016)].
- [2] P. D. Meerburg, D. Green, R. Flauger, et al., *Primordial Non-Gaussianity*, *Bull. Am. Astron. Soc.* **51** (May, 2019) 107, [[arXiv:1903.04409](https://arxiv.org/abs/1903.04409)].
- [3] A. Slosar, X. Chen, C. Dvorkin, et al., *Scratches from the Past: Inflationary Archaeology through Features in the Power Spectrum of Primordial Fluctuations*, *Bull. Am. Astron. Soc.* **51** (May, 2019) 98, [[arXiv:1903.09883](https://arxiv.org/abs/1903.09883)].
- [4] J. C. Hill, E. McDonough, M. W. Toomey, and S. Alexander, *Early Dark Energy Does Not Restore Cosmological Concordance*, *arXiv e-prints* (Mar., 2020) arXiv:2003.07355, [[arXiv:2003.07355](https://arxiv.org/abs/2003.07355)].
- [5] M. M. Ivanov, E. McDonough, J. C. Hill, et al., *Constraining Early Dark Energy with Large-Scale Structure*, *arXiv e-prints* (June, 2020) arXiv:2006.11235, [[arXiv:2006.11235](https://arxiv.org/abs/2006.11235)].

¹⁶https://github.com/lesgourg/class_public

¹⁷<https://github.com/sfschen/velocileptors>

¹⁸<https://hgomersall.github.io/pyFFTW/>

- [6] G. D’Amico, L. Senatore, P. Zhang, and H. Zheng, *The Hubble Tension in Light of the Full-Shape Analysis of Large-Scale Structure Data*, *arXiv e-prints* (June, 2020) arXiv:2006.12420, [[arXiv:2006.12420](#)].
- [7] A. Klypin, V. Poulin, F. Prada, et al., *Clustering and Halo Abundances in Early Dark Energy Cosmological Models*, *arXiv e-prints* (June, 2020) arXiv:2006.14910, [[arXiv:2006.14910](#)].
- [8] S.-F. Chen, Z. Vlah, and M. White, *Modeling features in the redshift-space halo power spectrum with perturbation theory*, *arXiv e-prints* (July, 2020) arXiv:2007.00704, [[arXiv:2007.00704](#)].
- [9] S. Ferraro et al., *Inflation and Dark Energy from Spectroscopy at $z > 2$* , [[arXiv:1903.09208](#)].
- [10] E. Castorina, S. Foreman, D. Karagiannis, et al., *Packed Ultra-wideband Mapping Array (PUMA): Astro2020 RFI Response*, *arXiv e-prints* (Feb., 2020) arXiv:2002.05072, [[arXiv:2002.05072](#)].
- [11] Planck Collaboration, Y. Akrami, F. Arroja, et al., *Planck 2018 results. X. Constraints on inflation*, *arXiv e-prints* (July, 2018) arXiv:1807.06211, [[arXiv:1807.06211](#)].
- [12] Planck Collaboration, P. A. R. Ade, N. Aghanim, et al., *Planck 2015 results. XX. Constraints on inflation*, *Astron. Astrophys.* **594** (Sept., 2016) A20, [[arXiv:1502.02114](#)].
- [13] A. Slosar, Z. Ahmed, D. Alonso, et al., *Packed Ultra-wideband Mapping Array (PUMA): A Radio Telescope for Cosmology and Transients*, in *Bull. Am. Astron. Soc.*, vol. 51, p. 53, Sep, 2019. [[arXiv:1907.12559](#)].
- [14] K. N. Abazajian, P. Adshead, Z. Ahmed, et al., *CMB- S_4 Science Book, First Edition*, *ArXiv e-prints* (Oct., 2016) [[arXiv:1610.02743](#)].
- [15] C. Dvorkin, M. Gerbino, D. Alonso, et al., *Neutrino Mass from Cosmology: Probing Physics Beyond the Standard Model*, *Bull. Am. Astron. Soc.* **51** (May, 2019) 64, [[arXiv:1903.03689](#)].
- [16] D. Green, M. A. Amin, J. Meyers, et al., *Messengers from the Early Universe: Cosmic Neutrinos and Other Light Relics*, *Bull. Am. Astron. Soc.* **51** (May, 2019) 159, [[arXiv:1903.04763](#)].
- [17] B. Jain and J. Khoury, *Cosmological tests of gravity*, *Annals of Physics* **325** (Jul, 2010) 1479–1516, [[arXiv:1004.3294](#)].
- [18] A. Joyce, L. Lombriser, and F. Schmidt, *Dark Energy Versus Modified Gravity*, *Annual Review of Nuclear and Particle Science* **66** (Oct., 2016) 95–122, [[arXiv:1601.06133](#)].
- [19] L. Amendola, S. Appleby, A. Avgoustidis, et al., *Cosmology and fundamental physics with the Euclid satellite*, *Living Reviews in Relativity* **21** (Apr., 2018) 2, [[arXiv:1606.00180](#)].
- [20] E. Castorina and M. White, *Measuring the growth of structure with intensity mapping surveys*, *Journal of Cosmology and Astro-Particle Physics* **2019** (June, 2019) 025, [[arXiv:1902.07147](#)].
- [21] Planck Collaboration, Y. Akrami, F. Arroja, et al., *Planck 2018 results. I. Overview and the cosmological legacy of Planck*, *arXiv e-prints* (July, 2018) arXiv:1807.06205, [[arXiv:1807.06205](#)].
- [22] S. Chabanier, M. Millea, and N. Palanque-Delabrouille, *Matter power spectrum: from Ly α forest to CMB scales*, *Mon. Not. Roy. Astron. Soc.* **489** (2019), no. 2 2247–2253, [[arXiv:1905.08103](#)].
- [23] DESI Collaboration, A. Aghamousa, J. Aguilar, et al., *The DESI Experiment Part I: Science, Targeting, and Survey Design*, *ArXiv e-prints* (Oct., 2016) [[arXiv:1611.00036](#)].
- [24] A. Blanchard, S. Camera, C. Carbone, et al., *Euclid preparation, a* **642** (Oct, 2020) A191, [[arXiv:1910.09273](#)].

- [25] S. Ilić et al., *Euclid preparation: XV. Forecasting cosmological constraints for the Euclid and CMB joint analysis*, [arXiv:2106.08346](#).
- [26] O. Doré, J. Bock, M. Ashby, et al., *Cosmology with the SPHEREX All-Sky Spectral Survey*, *arXiv e-prints* (Dec., 2014) [arXiv:1412.4872](#), [[arXiv:1412.4872](#)].
- [27] D. Schlegel, J. A. Kollmeier, and S. Ferraro, *The MegaMapper: a $z > 2$ spectroscopic instrument for the study of Inflation and Dark Energy*, in *Bull. Am. Astron. Soc.*, vol. 51, p. 229, Sept., 2019. [arXiv:1907.11171](#).
- [28] M. J. Wilson and M. White, *Cosmology with dropout selection: straw-man surveys & CMB lensing*, *Journal of Cosmology and Astro-Particle Physics* **2019** (Oct., 2019) 015, [[arXiv:1904.13378](#)].
- [29] The MSE Science Team, C. Babusiaux, M. Bergemann, et al., *The Detailed Science Case for the Maunakea Spectroscopic Explorer, 2019 edition*, *arXiv e-prints* (Apr., 2019) [arXiv:1904.04907](#), [[arXiv:1904.04907](#)].
- [30] W. J. Percival et al., *Cosmology with the MaunaKea Spectroscopic Explorer*, [arXiv:1903.03158](#).
- [31] R. Ellis and K. Dawson, *SpecTel: A 10-12 meter class Spectroscopic Survey Telescope*, in *Bull. Am. Astron. Soc.*, vol. 51, p. 45, Sept., 2019. [arXiv:1907.06797](#).
- [32] K. Dawson, J. Frieman, K. Heitmann, et al., *Cosmic Visions Dark Energy: Small Projects Portfolio*, [arXiv:1802.07216](#).
- [33] D. Spergel, N. Gehrels, C. Baltay, et al., *Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report*, *arXiv e-prints* (Mar., 2015) [arXiv:1503.03757](#), [[arXiv:1503.03757](#)].
- [34] K. Bandura, G. E. Addison, M. Amiri, et al., *Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder*, in *Ground-based and Airborne Telescopes V*, vol. 9145 of *Proc. SPIE*, p. 914522, July, 2014. [arXiv:1406.2288](#).
- [35] X. Chen, *The Tianlai Project: a 21CM Cosmology Experiment*, in *International Journal of Modern Physics Conference Series*, vol. 12 of *International Journal of Modern Physics Conference Series*, pp. 256–263, March, 2012. [arXiv:1212.6278](#).
- [36] K. Vanderlinde, A. Liu, B. Gaensler, et al., *The Canadian Hydrogen Observatory and Radio-transient Detector (CHORD)*, in *Canadian Long Range Plan for Astronomy and Astrophysics White Papers*, vol. 2020, p. 28, Oct., 2019. [arXiv:1911.01777](#).
- [37] L. B. Newburgh et al., *HIRAX: A Probe of Dark Energy and Radio Transients*, *Proc. SPIE Int. Soc. Opt. Eng.* **9906** (2016) 99065X, [[arXiv:1607.02059](#)].
- [38] F. B. Abdalla, P. Bull, S. Camera, et al., *Cosmology from HI galaxy surveys with the SKA, Advancing Astrophysics with the Square Kilometre Array (AASKA14)* (April, 2015) 17, [[arXiv:1501.04035](#)].
- [39] **Cosmic Visions 21 cm** Collaboration, R. Ansari et al., *Inflation and Early Dark Energy with a Stage II Hydrogen Intensity Mapping experiment*, [arXiv:1810.09572](#).
- [40] F. Villaescusa-Navarro, S. Genel, E. Castorina, et al., *Ingredients for 21 cm Intensity Mapping*, *Astrophys. J.* **866** (Oct, 2018) 135, [[arXiv:1804.09180](#)].
- [41] S.-F. Chen, E. Castorina, M. White, and A. Slosar, *Synergies between radio, optical and microwave observations at high redshift*, *Journal of Cosmology and Astro-Particle Physics* **2019** (July, 2019) 023, [[arXiv:1810.00911](#)].
- [42] R. K. Sheth, H. J. Mo, and G. Tormen, *Ellipsoidal collapse and an improved model for the number and spatial distribution of dark matter haloes*, *Monthly Notices of the Royal Astronomical Society* **323** (May, 2001) 1–12.

- [43] P. Bull, P. G. Ferreira, P. Patel, and M. G. Santos, *Late-time cosmology with 21cm intensity mapping experiments*, *Astrophys. J.* **803** (2015), no. 1 21, [[arXiv:1405.1452](#)].
- [44] H.-J. Seo, S. Dodelson, J. Marriner, et al., *A GROUND-BASED 21 cm BARYON ACOUSTIC OSCILLATION SURVEY*, *The Astrophysical Journal* **721** (aug, 2010) 164–173.
- [45] E. Castorina and F. Villaescusa-Navarro, *On the spatial distribution of neutral hydrogen in the universe: bias and shot-noise of the hI power spectrum*, *Monthly Notices of the Royal Astronomical Society* **471** (Jun, 2017) 1788–1796.
- [46] A. Obuljen, E. Castorina, F. Villaescusa-Navarro, and M. Viel, *High-redshift post-reionization cosmology with 21cm intensity mapping*, *JCAP* **05** (2018) 004, [[arXiv:1709.07893](#)].
- [47] A. Liu and J. R. Shaw, *Data Analysis for Precision 21 cm Cosmology*, *Publ. Astron. Soc. Pacific* **132** (June, 2020) 062001, [[arXiv:1907.08211](#)].
- [48] N. Galitzki, A. Ali, K. S. Arnold, et al., *The Simons Observatory: instrument overview, in Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, vol. 10708 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 1070804, July, 2018. [[arXiv:1808.04493](#)].
- [49] H. Sugai, P. A. R. Ade, Y. Akiba, et al., *Updated design of the cmb polarization experiment satellite litebird*, *Journal of Low Temperature Physics* **199** (Jan, 2020) 1107–1117.
- [50] J. Dunkley et al., *The Atacama Cosmology Telescope: likelihood for small-scale CMB data*, *JCAP* **07** (2013) 025, [[arXiv:1301.0776](#)].
- [51] F. McCarthy, J. C. Hill, and M. S. Madhavacheril, *Baryonic feedback biases on fundamental physics from lensed CMB power spectra*, [[arXiv:2103.05582](#)].
- [52] **NASA PICO** Collaboration, S. Hanany et al., *PICO: Probe of Inflation and Cosmic Origins*, [[arXiv:1902.10541](#)].
- [53] M. B. Bagley, C. Scarlata, V. Mehta, et al., *HST Grism-derived Forecasts for Future Galaxy Redshift Surveys*, *arXiv e-prints* (June, 2020) arXiv:2006.03602, [[arXiv:2006.03602](#)].
- [54] O. Doré, M. W. Werner, M. Ashby, et al., *Science Impacts of the SPHEREx All-Sky Optical to Near-Infrared Spectral Survey: Report of a Community Workshop Examining Extragalactic, Galactic, Stellar and Planetary Science*, *arXiv e-prints* (June, 2016) arXiv:1606.07039, [[arXiv:1606.07039](#)].
- [55] Planck Collaboration, N. Aghanim, Y. Akrami, et al., *Planck 2018 results. VI. Cosmological parameters*, *arXiv e-prints* (July, 2018) arXiv:1807.06209, [[arXiv:1807.06209](#)].
- [56] A. Font-Ribera, P. McDonald, N. Mostek, et al., *DESI and other dark energy experiments in the era of neutrino mass measurements*, *JCAP* **05** (2014) 023, [[arXiv:1308.4164](#)].
- [57] E. Schaan, S. Ferraro, and U. Seljak, *Photo-z outlier self-calibration in weak lensing surveys*, *JCAP* **12** (2020) 001, [[arXiv:2007.12795](#)].
- [58] B. Yu, R. Z. Knight, B. D. Sherwin, et al., *Towards Neutrino Mass from Cosmology without Optical Depth Information*, [[arXiv:1809.02120](#)].
- [59] M. Schmittfull, T. Baldauf, and M. Zaldarriaga, *Iterative initial condition reconstruction*, *Phys. Rev. D* **96** (July, 2017) 023505, [[arXiv:1704.06634](#)].
- [60] T. Brinckmann, D. C. Hooper, M. Archidiacono, et al., *The promising future of a robust cosmological neutrino mass measurement*, *JCAP* **01** (2019) 059, [[arXiv:1808.05955](#)].
- [61] S. Vagnozzi, T. Brinckmann, M. Archidiacono, et al., *Bias due to neutrinos must not uncorrect'd go*, *JCAP* **09** (2018) 001, [[arXiv:1807.04672](#)].
- [62] **LSST Dark Energy Science** Collaboration, D. Alonso et al., *The LSST Dark Energy Science Collaboration (DESC) Science Requirements Document*, [[arXiv:1809.01669](#)].

- [63] A. Chudaykin and M. M. Ivanov, *Measuring neutrino masses with large-scale structure: Euclid forecast with controlled theoretical error*, *Journal of Cosmology and Astro-Particle Physics* **2019** (Nov, 2019) 034, [[arXiv:1907.06666](#)].
- [64] A. Boyle and F. Schmidt, *Neutrino mass constraints beyond linear order: cosmology dependence and systematic biases*, *arXiv e-prints* (Nov., 2020) arXiv:2011.10594, [[arXiv:2011.10594](#)].
- [65] S.-F. Chen, H. Lee, and C. Dvorkin, *Precise and accurate cosmology with cmbxlss power spectra and bispectra*, 2021.
- [66] A. Mead and L. Verde, *Including beyond-linear halo bias in halo models*, *arXiv e-prints* (Nov., 2020) arXiv:2011.08858, [[arXiv:2011.08858](#)].
- [67] S.-F. Chen, Z. Vlah, E. Castorina, and M. White, *Redshift-Space Distortions in Lagrangian Perturbation Theory*, *arXiv e-prints* (Dec., 2020) arXiv:2012.04636, [[arXiv:2012.04636](#)].
- [68] D. Blas, J. Lesgourgues, and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes*, *Journal of Cosmology and Astro-Particle Physics* **7** (July, 2011) 034, [[arXiv:1104.2933](#)].
- [69] K. C. Chan, R. Scoccimarro, and R. K. Sheth, *Gravity and Large-Scale Non-local Bias*, *Phys. Rev. D* **85** (2012) 083509, [[arXiv:1201.3614](#)].
- [70] T. Baldauf, U. Seljak, V. Desjacques, and P. McDonald, *Evidence for quadratic tidal tensor bias from the halo bispectrum*, *Phys. Rev. D* **86** (Oct, 2012) 083540.
- [71] R. Takahashi, M. Sato, T. Nishimichi, et al., *Revising the Halofit Model for the Nonlinear Matter Power Spectrum*, *Astrophys. J.* **761** (Dec., 2012) 152, [[arXiv:1208.2701](#)].
- [72] F. Villaescusa-Navarro et al., *Ingredients for 21 cm Intensity Mapping*, *Astrophys. J.* **866** (2018), no. 2 135, [[arXiv:1804.09180](#)].
- [73] G. B. Field, *An Attempt to Observe Neutral Hydrogen Between the Galaxies.*, *Astrophys. J.* **129** (May, 1959) 525.
- [74] N. H. M. Crighton et al., *The neutral hydrogen cosmological mass density at $z = 5$* , *Mon. Not. Roy. Astron. Soc.* **452** (2015), no. 1 217–234, [[arXiv:1506.02037](#)].
- [75] F. Villaescusa-Navarro, F. Marulli, M. Viel, et al., *Cosmology with massive neutrinos I: towards a realistic modeling of the relation between matter, haloes and galaxies*, *JCAP* **03** (2014) 011, [[arXiv:1311.0866](#)].
- [76] E. Castorina, E. Sefusatti, R. K. Sheth, et al., *Cosmology with massive neutrinos II: on the universality of the halo mass function and bias*, *JCAP* **02** (2014) 049, [[arXiv:1311.1212](#)].
- [77] E. Castorina, C. Carbone, J. Bel, et al., *DEMNUi: The clustering of large-scale structures in the presence of massive neutrinos*, *JCAP* **07** (2015) 043, [[arXiv:1505.07148](#)].
- [78] F. Villaescusa-Navarro, A. Banerjee, N. Dalal, et al., *The imprint of neutrinos on clustering in redshift-space*, *Astrophys. J.* **861** (2018), no. 1 53, [[arXiv:1708.01154](#)].
- [79] M. LoVerde, *Halo bias in mixed dark matter cosmologies*, *Phys. Rev. D* **90** (2014), no. 8 083530, [[arXiv:1405.4855](#)].
- [80] J. B. Muñoz and C. Dvorkin, *Efficient Computation of Galaxy Bias with Neutrinos and Other Relics*, *Phys. Rev. D* **98** (2018), no. 4 043503, [[arXiv:1805.11623](#)].
- [81] C. Fidler, N. Sujata, and M. Archidiacono, *Relativistic bias in neutrino cosmologies*, *JCAP* **06** (2019) 035, [[arXiv:1812.09266](#)].
- [82] A. Aviles and A. Banerjee, *A Lagrangian Perturbation Theory in the presence of massive neutrinos*, *JCAP* **10** (2020) 034, [[arXiv:2007.06508](#)].

- [83] D. Wadekar, M. M. Ivanov, and R. Scoccimarro, *Cosmological constraints from BOSS with analytic covariance matrices*, *Phys. Rev. D* **102** (Dec., 2020) 123521, [[arXiv:2009.00622](#)].
- [84] A. LEWIS and A. CHALLINOR, *Weak gravitational lensing of the cmb*, *Physics Reports* **429** (Jun, 2006) 1–65.
- [85] M. LoVerde and N. Afshordi, *Extended limber approximation*, *Physical Review D* **78** (Dec, 2008).
- [86] M. Loverde and N. Afshordi, *Extended Limber approximation*, *Phys. Rev. D* **78** (Dec., 2008) 123506, [[arXiv:0809.5112](#)].
- [87] Krause, E. and Hirata, C. M., *Weak lensing power spectra for precision cosmology - multiple-deflection, reduced shear, and lensing bias corrections*, *A&A* **523** (2010) A28.
- [88] G. Pratten and A. Lewis, *Impact of post-Born lensing on the CMB*, *JCAP* **08** (2016) 047, [[arXiv:1605.05662](#)].
- [89] G. Marozzi, G. Fanizza, E. Di Dio, and R. Durrer, *CMB-lensing beyond the Born approximation*, *JCAP* **09** (2016) 028, [[arXiv:1605.08761](#)].
- [90] G. Fabbian, M. Calabrese, and C. Carbone, *CMB weak-lensing beyond the Born approximation: a numerical approach*, *JCAP* **02** (2018) 050, [[arXiv:1702.03317](#)].
- [91] G. Fabbian, A. Lewis, and D. Beck, *CMB lensing reconstruction biases in cross-correlation with large-scale structure probes*, *JCAP* **10** (2019) 057, [[arXiv:1906.08760](#)].
- [92] V. Böhm, C. Modi, and E. Castorina, *Lensing corrections on galaxy-lensing cross correlations and galaxy-galaxy auto correlations*, *JCAP* **03** (2020) 045, [[arXiv:1910.06722](#)].
- [93] C. Modi, M. White, and Z. Vlah, *Modeling cmb lensing cross correlations with cleft*, *Journal of Cosmology and Astroparticle Physics* **2017** (Aug, 2017) 009–009.
- [94] S. Foreman and L. Senatore, *The EFT of Large Scale Structures at All Redshifts: Analytical Predictions for Lensing*, *JCAP* **04** (2016) 033, [[arXiv:1503.01775](#)].
- [95] D. P. L. Bragança, M. Lewandowski, D. Sekera, et al., *Baryonic effects in the Effective Field Theory of Large-Scale Structure and an analytic recipe for lensing in CMB-S4*, [[arXiv:2010.02929](#)].
- [96] F. McCarthy, S. Foreman, and A. van Engelen, *Avoiding baryonic feedback effects on neutrino mass measurements from CMB lensing*, [[arXiv:2011.06582](#)].
- [97] E. Chung, S. Foreman, and A. van Engelen, *Baryonic effects on CMB lensing and neutrino mass constraints*, *Phys. Rev. D* **101** (2020), no. 6 063534, [[arXiv:1910.09565](#)]. [Erratum: *Phys.Rev.D* 102, 109903 (2020)].
- [98] E. Chung, S. Foreman, and A. van Engelen, *Baryonic effects on cmb lensing and neutrino mass constraints*, *Physical Review D* **101** (Mar, 2020).
- [99] C. Alcock and B. Paczynski, *An evolution free test for non-zero cosmological constant*, *Nature* **281** (Oct., 1979) 358.
- [100] H. Seo and D. J. Eisenstein, *Probing dark energy with baryonic acoustic oscillations from future large galaxy redshift surveys*, *The Astrophysical Journal* **598** (Dec, 2003) 720–740.
- [101] D. H. Weinberg, M. J. Mortonson, D. J. Eisenstein, et al., *Observational probes of cosmic acceleration*, *Phys. Rep.* **530** (Sept., 2013) 87–255, [[arXiv:1201.2434](#)].
- [102] D. J. Eisenstein, H.-J. Seo, E. Sirko, and D. N. Spergel, *Improving Cosmological Distance Measurements by Reconstruction of the Baryon Acoustic Peak*, *Astrophys. J.* **664** (Aug, 2007) 675–679, [[astro-ph/0604362](#)].

- [103] N. Padmanabhan, X. Xu, D. J. Eisenstein, et al., *A 2 per cent distance to $z = 0.35$ by reconstructing baryon acoustic oscillations - I. Methods and application to the Sloan Digital Sky Survey*, *Mon. Not. R. Astron. Soc.* **427** (Dec., 2012) 2132–2145, [[arXiv:1202.0090](#)].
- [104] S.-F. Chen, Z. Vlah, and M. White, *The reconstructed power spectrum in the Zeldovich approximation*, *Journal of Cosmology and Astro-Particle Physics* **2019** (Sept., 2019) 017, [[arXiv:1907.00043](#)].
- [105] S. Alam, M. Ata, S. Bailey, et al., *The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample*, *Mon. Not. R. Astron. Soc.* **470** (Sept., 2017) 2617–2652, [[arXiv:1607.03155](#)].
- [106] eBOSS Collaboration, S. Alam, M. Aubert, et al., *The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological Implications from two Decades of Spectroscopic Surveys at the Apache Point observatory*, *arXiv e-prints* (July, 2020) arXiv:2007.08991, [[arXiv:2007.08991](#)].
- [107] H. Seo and D. J. Eisenstein, *Improved forecasts for the baryon acoustic oscillations and cosmological distance scale*, *The Astrophysical Journal* **665** (Aug, 2007) 14–24.
- [108] A. E. Shapley, C. C. Steidel, M. Pettini, and K. L. Adelberger, *Rest-Frame Ultraviolet Spectra of $z \sim 3$ Lyman Break Galaxies*, *Astrophys. J.* **588** (May, 2003) 65–89, [[astro-ph/0301230](#)].
- [109] A. Verhamme, T. Garel, E. Ventou, et al., *Recovering the systemic redshift of galaxies from their Lyman alpha line profile*, *Mon. Not. R. Astron. Soc.* **478** (July, 2018) L60–L65, [[arXiv:1804.01883](#)].
- [110] A. Obuljen, F. Villaescusa-Navarro, E. Castorina, and M. Viel, *Baryon Acoustic Oscillations reconstruction with pixels*, *JCAP* **09** (2017) 012, [[arXiv:1610.05768](#)].
- [111] O. H. E. Philcox, M. M. Ivanov, M. Simonović, and M. Zaldarriaga, *Combining Full-Shape and BAO Analyses of Galaxy Power Spectra: A 1.6% CMB-independent constraint on H_0* , *JCAP* **05** (2020) 032, [[arXiv:2002.04035](#)].
- [112] C. Modi, M. White, A. Slosar, and E. Castorina, *Reconstructing large-scale structure with neutral hydrogen surveys*, *JCAP* **11** (2019) 023, [[arXiv:1907.02330](#)].
- [113] C. Modi, M. White, E. Castorina, and A. Slosar, *Mind the gap: the power of combining photometric surveys with intensity mapping*, [[arXiv:2102.08116](#)].
- [114] **Particle Data Group** Collaboration, P. A. Zyla et al., *Review of Particle Physics*, *PTEP* **2020** (2020), no. 8 083C01.
- [115] P. Collaboration, Y. Akrami, F. Arroja, et al., *Planck 2018 results. ix. constraints on primordial non-gaussianity*, 2019.
- [116] T. Essinger-Hileman, A. Ali, M. Amiri, et al., *CLASS: the cosmology large angular scale surveyor*, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII* (W. S. Holland and J. Zmuidzinas, eds.), vol. 9153 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 91531I, July, 2014. [[arXiv:1408.4788](#)].
- [117] S. Ferraro and K. M. Smith, *Characterizing the epoch of reionization with the small-scale CMB: Constraints on the optical depth and duration*, *Phys. Rev. D* **98** (2018), no. 12 123519, [[arXiv:1803.07036](#)].
- [118] M. A. Alvarez, S. Ferraro, J. C. Hill, et al., *Mitigating the optical depth degeneracy using the kinematic Sunyaev-Zel'dovich effect with CMB- S_4* , *Phys. Rev. D* **103** (2021), no. 6 063518, [[arXiv:2006.06594](#)].
- [119] A. Liu, J. R. Pritchard, R. Allison, et al., *Eliminating the optical depth nuisance from the CMB with 21 cm cosmology*, *Phys. Rev. D* **93** (2016), no. 4 043013, [[arXiv:1509.08463](#)].

- [120] T. S. Billings, P. L. Plante, and J. E. Aguirre, *Extracting the optical depth to reionization τ from 21 cm data using machine learning techniques*, 2021.
- [121] G. Mangano, G. Miele, S. Pastor, et al., *Relic neutrino decoupling including flavour oscillations*, *Nuclear Physics B* **729** (Nov, 2005) 221–234.
- [122] S. Bashinsky and U. Seljak, *Neutrino perturbations in CMB anisotropy and matter clustering*, *Phys. Rev. D* **69** (2004) 083002, [[astro-ph/0310198](#)].
- [123] D. Baumann, D. Green, J. Meyers, and B. Wallisch, *Phases of New Physics in the CMB*, *JCAP* **01** (2016) 007, [[arXiv:1508.06342](#)].
- [124] D. Baumann, D. Green, and M. Zaldarriaga, *Phases of New Physics in the BAO Spectrum*, *JCAP* **11** (2017) 007, [[arXiv:1703.00894](#)].
- [125] J. Maldacena, *Non-gaussian features of primordial fluctuations in single field inflationary models*, *Journal of High Energy Physics* **2003** (May, 2003) 013–013.
- [126] P. Creminelli and M. Zaldarriaga, *Single field consistency relation for the 3-point function*, *JCAP* **10** (2004) 006, [[astro-ph/0407059](#)].
- [127] F. Schmidt and M. Kamionkowski, *Halo clustering with nonlocal non-gaussianity*, *Physical Review D* **82** (Nov, 2010).
- [128] E. Castorina et al., *Redshift-weighted constraints on primordial non-Gaussianity from the clustering of the eBOSS DR14 quasars in Fourier space*, *JCAP* **09** (2019) 010, [[arXiv:1904.08859](#)].
- [129] N. Dalal, M. White, J. R. Bond, and A. Shirokov, *Halo Assembly Bias in Hierarchical Structure Formation*, *Astrophys. J.* **687** (Nov., 2008) 12–21, [[arXiv:0803.3453](#)].
- [130] A. Slosar, C. Hirata, U. Seljak, et al., *Constraints on local primordial non-Gaussianity from large scale structure*, *JCAP* **08** (2008) 031, [[arXiv:0805.3580](#)].
- [131] M. Biagetti, T. Lazeyras, T. Baldauf, et al., *Verifying the consistency relation for the scale-dependent bias from local primordial non-Gaussianity*, *Mon. Not. Roy. Astron. Soc.* **468** (2017), no. 3 3277–3288, [[arXiv:1611.04901](#)].
- [132] A. Barreira, *On the impact of galaxy bias uncertainties on primordial non-Gaussianity constraints*, *JCAP* **12** (2020) 031, [[arXiv:2009.06622](#)].
- [133] A. Moradinezhad Dizgah, M. Biagetti, E. Sefusatti, et al., *Primordial Non-Gaussianity from Biased Tracers: Likelihood Analysis of Real-Space Power Spectrum and Bispectrum*, [[arXiv:2010.14523](#)].
- [134] R. de Putter, *Primordial physics from large-scale structure beyond the power spectrum*, [[arXiv:1802.06762](#)].
- [135] D. Karagiannis, A. Lazanu, M. Liguori, et al., *Constraining primordial non-Gaussianity with bispectrum and power spectrum from upcoming optical and radio surveys*, *Mon. Not. Roy. Astron. Soc.* **478** (2018), no. 1 1341–1376, [[arXiv:1801.09280](#)].
- [136] J. Gleyzes, R. de Putter, D. Green, and O. Doré, *Biasing and the search for primordial non-gaussianity beyond the local type*, *Journal of Cosmology and Astroparticle Physics* **2017** (Apr, 2017) 002–002.
- [137] M. Schmittfull and U. Seljak, *Parameter constraints from cross-correlation of CMB lensing with galaxy clustering*, *Phys. Rev. D* **97** (June, 2018) 123540, [[arXiv:1710.09465](#)].
- [138] D. Baumann, *TASI Lectures on Inflation*, *arXiv e-prints* (July, 2009) arXiv:0907.5424, [[arXiv:0907.5424](#)].
- [139] J. Adams, B. Cresswell, and R. Easther, *Inflationary perturbations from a potential with a step*, *Physical Review D* **64** (Nov, 2001).

- [140] F. Beutler, M. Biagetti, D. Green, et al., *Primordial features from linear to nonlinear scales*, *Physical Review Research* **1** (Dec, 2019).
- [141] E. Linder, *The Rise of Dark Energy*, [arXiv:2106.xxxxx](#).
- [142] G. Efstathiou, *A Lockdown Perspective on the Hubble Tension (with comments from the SH0ES team)*, *arXiv e-prints* (July, 2020) arXiv:2007.10716, [[arXiv:2007.10716](#)].
- [143] V. Poulin, T. L. Smith, T. Karwal, and M. Kamionkowski, *Early dark energy can resolve the hubble tension*, *Physical Review Letters* **122** (Jun, 2019).
- [144] T. L. Smith, V. Poulin, J. L. Bernal, et al., *Early dark energy is not excluded by current large-scale structure data*, 2020.
- [145] T. Clifton, P. G. Ferreira, A. Padilla, and C. Skordis, *Modified gravity and cosmology*, *Physics Reports* **513** (Mar, 2012) 1–189.
- [146] C. M. Will, *The Confrontation between General Relativity and Experiment*, *Living Reviews in Relativity* **17** (Dec., 2014) 4, [[arXiv:1403.7377](#)].
- [147] A. Joyce, B. Jain, J. Khoury, and M. Trodden, *Beyond the cosmological standard model*, *Phys. Rep.* **568** (Mar, 2015) 1–98, [[arXiv:1407.0059](#)].
- [148] M. White, B. Reid, C.-H. Chuang, et al., *Tests of redshift-space distortions models in configuration space for the analysis of the BOSS final data release*, *Mon. Not. R. Astron. Soc.* **447** (Feb, 2015) 234–245, [[arXiv:1408.5435](#)].
- [149] S. Bharadwaj, *The Evolution of Correlation Functions in the Zeldovich Approximation and Its Implications for the Validity of Perturbation Theory*, *Astrophys. J.* **472** (Nov., 1996) 1–+, [[astro-ph/9](#)].
- [150] D. J. Eisenstein, H.-J. Seo, and M. White, *On the Robustness of the Acoustic Scale in the Low-Redshift Clustering of Matter*, *Astrophys. J.* **664** (Aug., 2007) 660–674, [[astro-ph/0604361](#)].
- [151] T. Matsubara, *Resumming cosmological perturbations via the Lagrangian picture: One-loop results in real space and in redshift space*, *Phys. Rev. D* **77** (Mar., 2008) 063530, [[arXiv:0711.2521](#)].
- [152] T. Matsubara, *Nonlinear perturbation theory with halo bias and redshift-space distortions via the Lagrangian picture*, *Phys. Rev. D* **78** (Oct., 2008) 083519, [[arXiv:0807.1733](#)].
- [153] M. Crocce and R. Scoccimarro, *Nonlinear evolution of baryon acoustic oscillations*, *Phys. Rev. D* **77** (Jan, 2008) 023533, [[arXiv:0704.2783](#)].
- [154] N. Padmanabhan, M. White, and J. D. Cohn, *Reconstructing baryon oscillations: A Lagrangian theory perspective*, *Phys. Rev. D* **79** (Mar., 2009) 063523, [[arXiv:0812.2905](#)].
- [155] Y. Noh, M. White, and N. Padmanabhan, *Reconstructing baryon oscillations*, *Phys. Rev. D* **80** (Dec., 2009) 123501, [[arXiv:0909.1802](#)].
- [156] J. Carlson, B. Reid, and M. White, *Convolution Lagrangian perturbation theory for biased tracers*, *Mon. Not. R. Astron. Soc.* **429** (Feb., 2013) 1674–1685, [[arXiv:1209.0780](#)].
- [157] S. Tassev and M. Zaldarriaga, *Towards an optimal reconstruction of baryon oscillations*, *Journal of Cosmology and Astro-Particle Physics* **10** (Oct., 2012) 006, [[arXiv:1203.6066](#)].
- [158] N. McCullagh and A. S. Szalay, *Nonlinear Behavior of Baryon Acoustic Oscillations from the Zel’dovich Approximation Using a Non-Fourier Perturbation Approach*, *Astrophys. J.* **752** (June, 2012) 21, [[arXiv:1202.1306](#)].
- [159] M. White, *The Zel’dovich approximation*, *Mon. Not. R. Astron. Soc.* **439** (Apr., 2014) 3630–3640, [[arXiv:1401.5466](#)].

- [160] L. Senatore and M. Zaldarriaga, *The IR-resummed Effective Field Theory of Large Scale Structures*, *Journal of Cosmology and Astro-Particle Physics* **2** (Feb., 2015) 13, [[arXiv:1404.5954](#)].
- [161] M. Schmittfull, Y. Feng, F. Beutler, et al., *Eulerian BAO reconstructions and N -point statistics*, *Phys. Rev. D* **92** (Dec., 2015) 123522, [[arXiv:1508.06972](#)].
- [162] T. Baldauf, M. Mirbabayi, M. Simonović, and M. Zaldarriaga, *Equivalence principle and the baryon acoustic peak*, *Phys. Rev. D* **92** (Aug., 2015) 043514, [[arXiv:1504.04366](#)].
- [163] Z. Vlah, U. Seljak, M. Yat Chu, and Y. Feng, *Perturbation theory, effective field theory, and oscillations in the power spectrum*, *Journal of Cosmology and Astro-Particle Physics* **2016** (Mar, 2016) 057, [[arXiv:1509.02120](#)].
- [164] M. McQuinn and M. White, *Cosmological perturbation theory in 1+1 dimensions*, *Journal of Cosmology and Astro-Particle Physics* **1** (Jan., 2016) 043, [[arXiv:1502.07389](#)].
- [165] D. Blas, M. Garny, M. M. Ivanov, and S. Sibiryakov, *Time-sliced perturbation theory II: baryon acoustic oscillations and infrared resummation*, *Journal of Cosmology and Astro-Particle Physics* **2016** (July, 2016) 028, [[arXiv:1605.02149](#)].
- [166] H.-J. Seo, F. Beutler, A. J. Ross, and S. Saito, *Modeling the reconstructed BAO in Fourier space*, *Mon. Not. R. Astron. Soc.* **460** (Aug., 2016) 2453–2471, [[arXiv:1511.00663](#)].